

---

# KNIME4NGS: a comprehensive toolbox for high-throughput Next Generation Sequencing analysis - User manual

---

Maximilian Hastreiter, Tim Jeske, Jonathan Hoser, Michael Kluge, Kaarin Ahomaa,  
Marie-Sophie Friedl, Sebastian J. Kopetzky, Jan-Dominik Quell,  
H.-Werner Mewes and Robert Küffner  
December 8, 2016

## CONTENTS

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Prerequisites . . . . .	3
1.2	Installation . . . . .	3
<b>2</b>	<b>Basic Usage</b>	<b>5</b>
2.1	File Loading and Pipeline Entry Point . . . . .	5
2.2	Read Filtering and Quality Control . . . . .	5
2.3	Mapping and SAM-BAM Conversion . . . . .	5
2.4	Alignment Refinement . . . . .	6
2.5	Variant Calling and Filtering . . . . .	6
2.6	Variant Annotation . . . . .	6
2.7	Differential Expression . . . . .	7
<b>3</b>	<b>Parallel Execution</b>	<b>8</b>
3.1	Installation . . . . .	8
3.2	Usage . . . . .	8
<b>4</b>	<b>The HTE</b>	<b>8</b>
4.1	HTE Database . . . . .	9
4.2	Klock Files . . . . .	9
<b>5</b>	<b>KNIME Cluster Executor</b>	<b>9</b>
<b>6</b>	<b>Preference Page</b>	<b>10</b>
<b>7</b>	<b>FAQ</b>	<b>12</b>



# 1 INTRODUCTION

Knime4NGS is a set of custom KNIME nodes, designed for processing large next generation sequencing datasets. The nodes wrap functionalities of popular command line tools like BWA, Samtools and GATK and allow an easy creation of comprehensive NGS KNIME workflows.

## 1.1 PREREQUISITES

- Linux System (required)
- KNIME or KNIME SDK (required)
- KNIME Cluster Extension<sup>1</sup> (recommended)
- KNIME Virtual Nodes Extension (for parallel execution only, see section 3))
- R installation with packages: gplots, ggplot2, reshape, gsalib; for RNA-Seq pipeline: argparse, edgeR, DESeq, Biobase and limma
- Genome Analysis Toolkit (GATK) 3.x <sup>2</sup> (tested with GATK 3.5 and 3.6)
- Variant Effect Predictor (VEP) script including the Ensembl Core and Variation APIs (for variant annotation, see section 2.6)<sup>3</sup>. VEP itself requires the following perl modules: Archive::Extract, Archive::Zip, Test::More, DBI and CGI

## 1.2 INSTALLATION

1. Start KNIME
2. Specify the KNIME4NGS Update Site as a software repository within KNIME. In order to do that, go to Help -> Install New Software... and add `http://ibisngs.github.io/knime4ngs/updateSite` to the available software sites as shown in Figure 1.1.
3. Select the new software site. The IBIS KNIME Nodes should now appear. Select and install the displayed package. After restarting KNIME, the KNIME4NGS nodes should appear in the KNIME Node Repository as shown in Figure 1.2.
4. (Optional) Configure the KNIME4NGS Preference Page which can be found at: File -> Preferences -> KNIME -> KNIME4NGS.
  - Missing binaries can be obtained by using 'Download missing binaries' and/or search for existing binaries in your file system. This step is necessary since most nodes do not contain the respective execution binaries. Although binaries can be selected in each node independently, setting global binary paths in the preference page is recommended. See section 6 for more details.
  - Activate HTE execution. See section 4 for more details.
5. (Optional) Go to the KNIME GUI Preference Page and set Console View Log Level to DEBUG to obtain more logging information.

---

<sup>1</sup><https://www.knime.org/cluster-execution>

<sup>2</sup><https://www.broadinstitute.org/gatk/>

<sup>3</sup>[http://www.ensembl.org/info/docs/tools/vep/script/vep\\_download.html](http://www.ensembl.org/info/docs/tools/vep/script/vep_download.html)

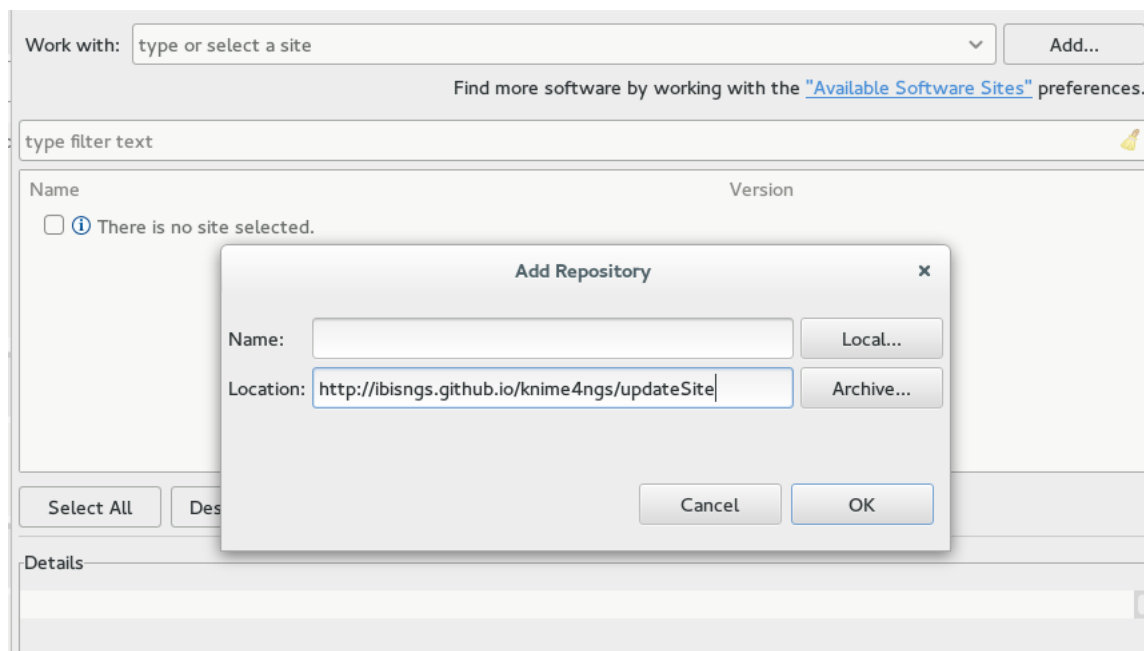


Figure 1.1: Add KNIME4NGS Update Site as a repository site.




Figure 1.2: KNIME4NGS nodes appear in the KNIME Node Repository.

## 2 BASIC USAGE

### 2.1 FILE LOADING AND PIPELINE ENTRY POINT

As an entry point to your pipeline, you can and should use the **FileLoader** node for selecting your input files. Depending on your input, this node can be used to enter pipelines at any stage. Possible inputs include FastA, FastQ, SAM/BAM or VCF/gVCF files. After node execution, the selected file path(s) will be displayed in the nodes output table.




Figure 2.1: Loading files using the FileLoader node.

### 2.2 READ FILTERING AND QUALITY CONTROL

For basic quality control and read filtering, we have included a modified version of **FastQC**, which is designed to work with our proposed read-filtering and trimming tool, the **RawReadManipulator**. While the FastQC node will prepare summary statistics about your input sequences, it also provides suggested filtering criteria depending on the input quality. These filter settings are then used by the **RawReadManipulator** to filter and trim the data. Alternatively, you can select your own settings in the node dialog.




Figure 2.2: The RawReadManipulator can use proposed filter settings created by FastQC.

A suggested workflow is depicted in Figure 2.3. Running **FastQC** for a second time after read filtering allows you to check the quality of the filtered FastQ files.

### 2.3 MAPPING AND SAM-BAM CONVERSION

For mapping the reads against a reference genome, you can choose between several mapping nodes like **BWA**, **Segemehl** or **Star**, depending on your type of analysis. If you have already used the previous nodes for quality control, you can directly connect them to your preferred mapping node, as shown in Figure 2.4. Otherwise use the **FileLoader** to load your file(s). SAM files can be converted to BAM format using the **PicardTools** node.



Figure 2.3: Quality Check and Filtering Workflow

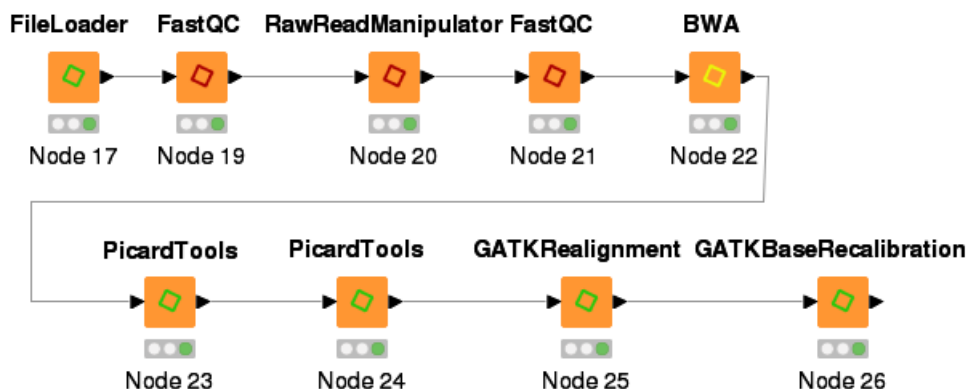


Figure 2.4: Mapping, Conversion and Alignment Refinement

## 2.4 ALIGNMENT REFINEMENT

In order to improve alignment quality and therefore the quality of all following analysis steps, there are several ways to increase alignment reliability. We have included several nodes, following the GATK Base Practices, for marking PCR duplicates (**PicardTools**), Indel Realignment (**GATKRealignment**) and Base Quality Score Recalibration (**GATKBaseRecalibration**) as shown in Figure 2.4. Note, that the latter two GATK nodes require read groups. If there are no read groups in your BAM files you can add some by using the *AddOrReplaceReadGroups* tool of the **PicardTools** node.

## 2.5 VARIANT CALLING AND FILTERING

Variant Calling can be done in different ways, depending on the field of interest. For SNP and short Indel calling, we provide the **UnifiedGenotyper** and **HaplotypeCaller** (gvcf mode only). The resulting GVCf files of the **HaplotypeCaller** can be converted to VCF files by using **GenotypeGVCFs**. Furthermore, we included **Pindel**, an algorithm for calling long Indels and structural variants (currently, the node only supports single-sample calling). **Pindel** requires a config file which can easily be generated by the *CollectInsertMetrics* tool of **PicardTools**. All mentioned variant calling nodes are shown in Figure 2.5.

To exclude low quality variants that arise from sequencing artefacts we suggest to apply Variant Quality Score Recalibration (**VQSR**) based on GATK's VariantRecalibration and ApplyRecalibration tool. Additionally, the **VCFToolsFilter** enables filtering for further sequencing quality parameters.

## 2.6 VARIANT ANNOTATION

Our toolbox provides the **VEP** node to annotate the functional impact of called variants. The node wraps Ensembl's VEP script and allows you to set the most important flags directly via the node dialog. Nevertheless, it is possible to specify further flags or plugins in an extra text




Figure 2.5: For Variant Calling, three different nodes are provided, each based on different calling algorithms.

field. After applying **VEP**, the variants can be filtered according to the annotations by applying the **VEPFilter**.

## 2.7 DIFFERENTIAL EXPRESSION

For differential expression analysis, we provide a pipeline based on the Subread package **featureCounts**. It can be used for generating raw read counts that serve as input for typical differential expression analysis like **DESeq**, **Limma** or **EdgeR**. An exemplary workflow is shown in Figure 2.6.

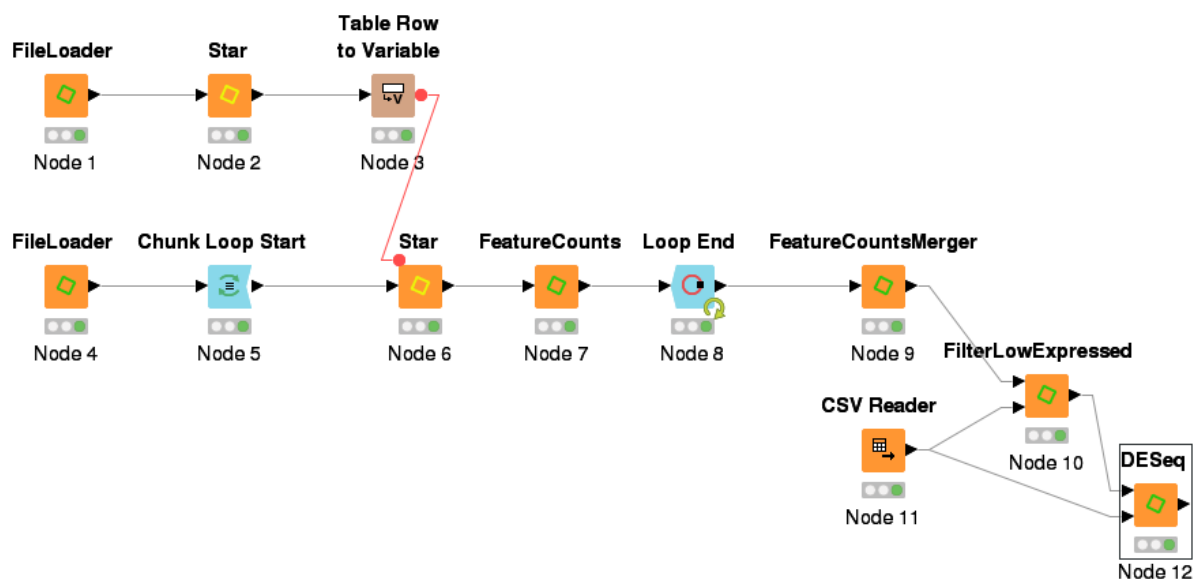


Figure 2.6: Differential Expression Workflow using Star and DESeq

## 3 PARALLEL EXECUTION

Analysing a large number of samples typically requires parallel processing to ensure reasonable runtimes. In KNIME, this can be achieved by using the **Parallel Chunk** nodes included in the **KNIME Virtual Nodes Extension**.

### 3.1 INSTALLATION

In order to install the extension, select *Help -> Install New Software...* Choose the official KNIME Update Site. Select and install the KNIME Virtual Nodes Extension. After restarting KNIME, the Parallel Chunk nodes should be available in the node repository (KNIME Labs category).

### 3.2 USAGE

As shown in Figure 3.1, these nodes can be combined with the KNIME4NGS nodes to create a parallel execution pipeline for large scale NGS analysis.

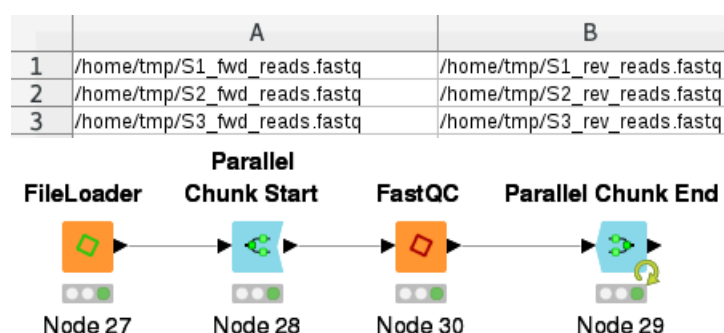


Figure 3.1: Exemplary parallel execution workflow

To run a pipeline in parallel mode, first prepare a file containing all input filepaths (e.g. Col0: forward read files, Col1: reverse read files, one row for each sample). Import the files list using the **FileLoader** and check the box: "Handle input file as list of file paths?". By using the **Parallel Chunk Start**, each row will be handled separately if the chunk count equals the number of rows. All nodes located between the Parallel Chunk Start and End will then be executed in parallel.

## 4 THE HTE

Several nodes of our toolbox can be run in the High Throughput Execution (HTE) mode which has three main benefits. Firstly, nodes can automatically be restarted if they fail for some reason. Secondly, the occurring errors are reported in a database and you are notified by Email if your workflow fails. Thirdly, HTE remembers the successful execution of a node which prevents the time-consuming re-execution.

In some cases, tools used by the KNIME4NGS nodes fail because of errors that could be solved by restarting the tool. HTE is capable of restarting nodes automatically as often as you wish. You can globally set a threshold for the number of restarts for all nodes in the preference page (see section 6) or set it individually for each node in the configuration dialog. This also prevents restarting tools that continuously fail. In this case, manual inspection by the user is required. Therefore, we provide a HTE-NodeView for stdError and stdout log messages that includes relevant information for debugging. The HTE-NodeView can be accessed by right-clicking on the respective node. Additional log files can also be found at the output file location.



## 4.1 HTE DATABASE

If you use the HTE mode you can create a HTE database which logs errors occurring during the execution of your workflows. The database is stored locally as a file and requires no database server. The HTE database consists of two tables: the HTEExecution table comprises entries for each executed node and the HTEError\_History table contains the error messages if the execution failed at least once. If you want to use a HTE database in your KNIME installation, go to the KNIME4NGS preference page and create a new database file (hte.db) or select an existing one.

You can easily access the database using two KNIME nodes that are included in the basic repository. An exemplary workflow is shown in Figure 4.1. Simply add the path to your HTE database (the hte.db file) in the configuration dialog of the SQLite Connector node. The subsequent Database Reader allows you to inspect the fields of both tables via "Fetch Metadata" and to execute SQL queries.




Figure 4.1: Read the HTE database

## 4.2 KLOCK FILES

Sometimes, it may happen that you reset a node or a complete workflow unintentionally. In general, nodes and the respective workflows are completely re-executed after resetting. This behavior can cause you to wait for a long time to get the same results although you have not changed the configuration of a node. For that reason, most of our nodes generate klock files. Within a klock file, the execution command, host and termination state are stored. Prior to the (re-)execution, all nodes check for the presence of klock files in order to determine if an execution is actually required. This is realised by comparing the execution command strings and checking the previous termination state. Therefore, these files can prevent the re-execution of successfully executed nodes if their configuration (execution command string) is not changed. If you want to re-execute a node without altering its configuration you can simply delete the corresponding klock file which is located in the same directory as the output file.

## 5 KNIME CLUSTER EXECUTOR

The KNIME Cluster Executor provides a direct connection between KNIME and a high-performance computing cluster. This allows to distribute node execution across several computation hosts which can significantly increase the processing speed of KNIME workflows (For more information about the KNIME Cluster Executor and how to obtain it, see <https://www.knime.org/knime-cluster-executor>).

We highly recommend using KNIME4NGS in combination with the KNIME Cluster Executor for all NGS applications that go beyond processing data sets with more than a few samples (depending on actual file size). Nevertheless, also single high-powered workstations can be sufficient for processing larger datasets. In this case, no Cluster Executor is required as the jobs can be distributed across the single machine by using the Parallel Chunk Environment.

In general, it always depends on the actual data size, number of samples as well as the available workstations and time in order to determine the best solution.

## 6 PREFERENCE PAGE

The toolbox comes along with its own binary management system controlled via the preference page shown in Figure 6.1. This page allows you to set several options which simplifies and accelerates the configuration of your workflows. All nodes will also work without the preference page settings but they are much more comfortable to use if you take the time to configure the preference page.




Figure 6.1: KNIME4NGS Preference Page. It can be found in "File > Preferences > KNIME > KNIME4NGS"

The first part of the preference page offers you three possibilities to set paths to binaries required by the nodes of the toolkit. First, you can specify a directory where missing binaries will be searched. Please consider that also subdirectories are included and thus the searching process may take a while depending on the selected starting directory. Second, it is also possible to download missing binaries in your directory of choice. Executable binaries will be saved into the selected folder and automatically integrated into the preference page without the need for further configuration. Note, that not all binaries can be download due to license issues. If required, you can set each binary path manually.

In the middle part of the page you can set the reference genome. As all GATK and some

other nodes require a reference genome, setting it only once will save you a lot of time.

The last part of the page allows you to set HTE preferences (see also section 4). If you decide to use HTE, you have to create a new HTE database file (hte.db) or you can select an existing database file. Further, you can set the threshold for the repeated execution in the preference page which applies to all nodes. If you want to get an Email whenever a node fails after repeated execution, you have to set an Email host and an address to send and receive Emails.

## 7 FAQ

1. **Which nodes can work on compressed input files?**
  - Compatible with compressed FastQ files: FastQC, RawReadManipulator
  - Compatible with compressed vcf files: VEP
2. **No active session error in parallel chunk environment.**
  - Reset parallel chunk start and re-execute.
3. **The node wont execute and always returns to "orange configure-state" without error message.**
  - Reset the last successfully executed predecessor node and re-execute the nodes.
4. **Update has encountered a problem.**
  - Restart KNIME before updating.
5. **Node is in EXECUTING state and cannot be canceled although it has been successfully executed on the cluster.**
  - Go to the corresponding workflow folder in the knime-workspace. Edit the xml file of the executing node and replace EXECUTING by IDLE or CONFIGURED.

## GATK ERRORS

1. **VQSR: Unable to retrieve results.**
  - Use VQSR in single-threaded mode (Threads 1).
2. **VQSR: No data found.**
  - Reduce maximum number of Gaussians to 4.

## 8 OVERVIEW OF ALL KNIME4NGS NODES

Node	Description	Downloadable binaries (version)
<b>Bcftools</b>	A set of utilities that manipulate variant calls in the Variant Call Format (VCF) and its binary counterpart BCF. [14]	✓ (version 1.3)
<b>Bowtie2</b>	An ultrafast and memory-efficient tool for aligning sequencing reads to long reference sequences. [11]	✓ (version 2.2.7)
<b>BWA</b>	The Burrows-Wheeler Alignment Tool (BWA) aligns relatively short sequences (queries) to a sequence database (target), such as the human reference genome. [12],[13]	✓ (version 0.7.12)
<b>DepthOfCoverage</b>	Assess sequence coverage by a wide array of metrics, partitioned by sample, read group, or library. [17],[6],[21]	
<b>DESeq</b>	This node uses the DESeq package to test for differential expression based on a model using the negative binomial distribution. The node tries to automatically install the missing R packages during execution. If this fails, we recommend to manually install the packages via the Bioconductor. [1]	
<b>EdgeR</b>	Assessing differential expression in comparative RNA-seq experiments. The node tries to automatically install the missing R packages during execution. If this fails, we recommend to manually install the packages via the Bioconductor. [20]	
<b>FastQC</b>	Quality check for NGS read data. [2]	
<b>FastqcStatisticMerger</b>	Merges statistics of FastQC modules.	
<b>FeatureCounts</b>	An efficient general purpose program for assigning sequence reads to genomic features. [16]	✓ (version 1.4.4)
<b>FeatureCountsMerger</b>	Merges the results of multiple feature count runs.	
<b>FeatureCountsStatisticMerger</b>	Merges statistics of the FeatureCounts node.	
<b>FileLoader</b>	Loads either one or two separate input file(s) or a list of files whereby it serves as a starting point for any NGS workflow.	
<b>FilterLowExpressed</b>	Filters low expressed genes before differential expression analysis. The node tries to automatically install the missing R packages during execution. If this fails, we recommend to manually install the packages via the Bioconductor.	
<b>GATKBaseRecalibration</b>	Recalibration of sequencing quality scores. [17],[6],[21]	
<b>CombineGVCFs</b>	Combines any number of gVCF files that were produced by the Haplotype Caller into a single joint gVCF file. [17],[6],[21]	
<b>CombineVCFs</b>	Combines any number of VCF lists that were produced by Pindel into a single joint VCF file.	

<b>GenotypeGVCFs</b>	Performs joint genotyping on gVCF files produced by HaplotypeCaller. [17],[6],[21]	
<b>GATKHaplotypeCaller</b>	Calls germline SNPs and indels via local re-assembly of haplotypes. [17],[6],[21]	
<b>MergeTwoVCFs</b>	Merges two VCF files into one combined VCF file.	
<b>GATKPhaseByTransmission</b>	Computes the most likely genotype combination and phases trios and parent/child pairs. [17],[6],[21]	
<b>GATKRealignment</b>	Realignment around insertions and deletions. [17],[6],[21]	
<b>GATKSelectVariants</b>	Selects a subset of variants from a larger callset. [17],[6],[21]	
<b>GATKUnifiedGenotyper</b>	Calls SNPs and indels. [17],[6],[21]	
<b>GATKVariantFiltration</b>	Filters variant calls based on INFO and FORMAT annotations. [17],[6],[21]	
<b>KGGSeq</b>	Filters and prioritizes genetic variants from whole exome sequencing data. [15]	
<b>Limma</b>	Detects differential expressed genes by using the limma package. The node tries to automatically install the missing R packages during execution. If this fails, we recommend to manually install the packages via the Bioconductor. [19]	
<b>PicardTools</b>	Tools to manipulate SAM/BAM files. [10]	
<b>Pindel</b>	A pattern growth approach to detect breakpoints of large deletions and medium sized insertions from paired-end reads. [22]	✓ (version 0.2.5a8)
<b>RawReadManipulator</b>	Filters reads from fastQ/fastA files.	
<b>RawReadManipulatorStatisticMerger</b>	Merges statistics of the RawReadManipulator node.	
<b>SamTools</b>	Provides access to many functionalities of SAM-tools. [14]	✓ (version 1.1)
<b>Segemehl</b>	Heuristic mapping of short read sequences with gaps. [9],[8]	✓ (version 0.2.0-418)
<b>SnEff</b>	Annotates and predicts the effects of genetic variants. [4]	
<b>SnSift</b>	A wrapper for the SnSift toolbox. [3]	
<b>Star</b>	Aligns RNA-seq reads to a reference genome using uncompressed suffix arrays and is "splice aware". [7]	✓
<b>StarStatisticMerger</b>	Merges statistics of the STAR node.	
<b>VCFMerger</b>	Combines VCF records from different sources.	
<b>VCFSorter</b>	Sorts the chromosomes of an input VCF file analogous to a given reference sequence.	
<b>VCFToolsFilter</b>	Filters annotated VCF files. [5]	
<b>VEP</b>	The Variant Effect Predictor (VEP) annotates genomic variants. [18]	
<b>VEPFilter</b>	Filters VCF files annotated by VEP.	

---

<b>VQSR</b>	Variant Quality Score Recalibration: creates a Gaussian mixture model by looking at the annotations values over a high quality subset of the input call set and then evaluates all input variants. [17],[6],[21]
-------------	--

---

Table 8.1: Overview of all KNIME4NGS nodes

## REFERENCES

- [1] S. Anders and W. Huber. Differential expression analysis for sequence count data. *Genome Biology*, 11:R106, 2010.
- [2] S. Andrews. Babraham Bioinformatics - FastQC A Quality Control tool for High Throughput Sequence Data, 2010.
- [3] P. Cingolani, V. M. Patel, M. Coon, T. Nguyen, S. J. Land, D. M. Ruden, and X. Lu. Using *Drosophila melanogaster* as a Model for Genotoxic Chemical Mutational Studies with a New Program, SnpSift. *Frontiers in Genetics*, 3:35, 2012.
- [4] P. Cingolani, A. Platts, L. L. Wang, M. Coon, T. Nguyen, L. Wang, S. J. Land, X. Lu, and D. M. Ruden. A program for annotating and predicting the effects of single nucleotide polymorphisms, SnpEff: SNPs in the genome of *Drosophila melanogaster* strain w1118; iso-2; iso-3. *Fly*, 6(2):80–92, June 2012.
- [5] P. Danecek, A. Auton, G. Abecasis, C. A. Albers, E. Banks, M. A. DePristo, R. E. Handsaker, G. Lunter, G. T. Marth, S. T. Sherry, G. McVean, R. Durbin, and 1000 Genomes Project Analysis Group. The variant call format and VCFtools. *Bioinformatics*, 27(15):2156–2158, Aug. 2011.
- [6] M. A. DePristo, E. Banks, R. Poplin, K. V. Garimella, J. R. Maguire, C. Hartl, A. A. Philippakis, G. del Angel, M. A. Rivas, M. Hanna, A. McKenna, T. J. Fennell, A. M. Kernytzsky, A. Y. Sivachenko, K. Cibulskis, S. B. Gabriel, D. Altshuler, and M. J. Daly. A framework for variation discovery and genotyping using next-generation DNA sequencing data. *Nature Genetics*, 43(5):491–498, May 2011.
- [7] A. Dobin, C. A. Davis, F. Schlesinger, J. Drenkow, C. Zaleski, S. Jha, P. Batut, M. Chaisson, and T. R. Gingeras. STAR: ultrafast universal RNA-seq aligner. *Bioinformatics*, 29(1):15–21, Jan. 2013.
- [8] S. Hoffmann, C. Otto, G. Doose, A. Tanzer, D. Langenberger, S. Christ, M. Kunz, L. M. Holdt, D. Teupser, J. Hackermüller, and P. F. Stadler. A multi-split mapping algorithm for circular RNA, splicing, trans-splicing and fusion detection. *Genome Biology*, 15(2):R34, 2014.
- [9] S. Hoffmann, C. Otto, S. Kurtz, C. M. Sharma, P. Khaitovich, J. Vogel, P. F. Stadler, and J. Hackermüller. Fast Mapping of Short Sequences with Mismatches, Insertions and Deletions Using Index Structures. *PLoS Computational Biology*, 5(9):e1000502, Sept. 2009.
- [10] B. Institute. Picard Tools - By Broad Institute, 2009.
- [11] B. Langmead and S. L. Salzberg. Fast gapped-read alignment with Bowtie 2. *Nature Methods*, 9(4):357–359, Mar. 2012.
- [12] H. Li and R. Durbin. Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics (Oxford, England)*, 25(14):1754–1760, July 2009.
- [13] H. Li and R. Durbin. Fast and accurate long-read alignment with Burrows-Wheeler transform. *Bioinformatics (Oxford, England)*, 26(5):589–595, Mar. 2010.
- [14] H. Li, B. Handsaker, A. Wysoker, T. Fennell, J. Ruan, N. Homer, G. Marth, G. Abecasis, R. Durbin, and 1000 Genome Project Data Processing Subgroup. The Sequence Alignment/Map format and SAMtools. *Bioinformatics*, 25(16):2078–2079, Aug. 2009.
- [15] M.-X. Li, H.-S. Gui, J. S. H. Kwan, S.-Y. Bao, and P. C. Sham. A comprehensive framework for prioritizing variants in exome sequencing studies of Mendelian diseases. *Nucleic Acids Research*, 40(7):e53, Apr. 2012.

- [16] Y. Liao, G. K. Smyth, and W. Shi. featureCounts: an efficient general purpose program for assigning sequence reads to genomic features. *Bioinformatics (Oxford, England)*, 30(7):923–930, Apr. 2014.
- [17] A. McKenna, M. Hanna, E. Banks, A. Sivachenko, K. Cibulskis, A. Kernytsky, K. Garimella, D. Altshuler, S. Gabriel, M. Daly, and M. A. DePristo. The Genome Analysis Toolkit: a MapReduce framework for analyzing next-generation DNA sequencing data. *Genome Research*, 20(9):1297–1303, Sept. 2010.
- [18] W. McLaren, L. Gil, S. E. Hunt, H. S. Riat, G. R. S. Ritchie, A. Thormann, P. Flicek, and F. Cunningham. The Ensembl Variant Effect Predictor. *Genome Biology*, 17(1), Dec. 2016.
- [19] M. E. Ritchie, B. Phipson, D. Wu, Y. Hu, C. W. Law, W. Shi, and G. K. Smyth. limma powers differential expression analyses for RNA-sequencing and microarray studies. *Nucleic Acids Research*, 43(7):e47, Apr. 2015.
- [20] M. D. Robinson, D. J. McCarthy, and G. K. Smyth. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics (Oxford, England)*, 26(1):139–140, Jan. 2010.
- [21] G. A. Van der Auwera, M. O. Carneiro, C. Hartl, R. Poplin, G. Del Angel, A. Levy-Moonshine, T. Jordan, K. Shakir, D. Roazen, J. Thibault, E. Banks, K. V. Garimella, D. Altshuler, S. Gabriel, and M. A. DePristo. From FastQ data to high confidence variant calls: the Genome Analysis Toolkit best practices pipeline. *Current Protocols in Bioinformatics*, 43:11.10.1–33, 2013.
- [22] K. Ye, M. H. Schulz, Q. Long, R. Apweiler, and Z. Ning. Pindel: a pattern growth approach to detect break points of large deletions and medium sized insertions from paired-end short reads. *Bioinformatics*, 25(21):2865–2871, Nov. 2009.