

Introduction to Sentiment Analysis

-- *Session 3: Learning Subjectivity*--

Winter Semester 2019/2020

Instructor: Michael Wiegand

Institute for Computational Linguistics

Heidelberg University, Germany

Outline

- ▶ Learning Subjective Adjectives
- ▶ Learning Subjective Nouns using Extraction Pattern Bootstrapping

Outline

- ▶ **Learning Subjective Adjectives**
- ▶ Learning Subjective Nouns using Extraction Pattern Bootstrapping

“Learning Subjective Adjectives”

By Janyce M. Wiebe,

In *Proceedings of AAAI*, 2000.

Subjective vs. Objective

- ▶ Text/sentence level:
 - ▶ subjective:
 - ▶ *At several different layers, it's a fascinating tale.*
 - ▶ objective:
 - ▶ *Bell Industries Inc. increased its quarterly to 10 cents from 7 cents a share.*
- ▶ Word level (i.e. adjectives):
 - ▶ subjective:
 - ▶ *nice, bad, huge, tall, possibly*
 - ▶ objective:
 - ▶ *round, American, financial, plastic*

Definition of Subjectivity

- ▶ Focus on 3 types of subjectivity:
 - ▶ positive evaluation: *interesting*, *awesome*
 - ▶ negative evaluation: *horrible*, *bad*
 - ▶ speculation: *probably*, *likely*

Why Learning Subjective Adjectives?

- ▶ Bruce & Wiebe (2000): probability that a sentence is subjective, given that there is at least 1 adjective, is 55.8%.
- ▶ Classifying all sentences as subjective provides us with a much lower precision.
- ▶ Conclusion:
 - ▶ Adjectives, in general, are a good indicator of subjective language.
 - ▶ But not all adjectives are subjective.

Why Learning Subjective Adjectives?

- ▶ Bruce & Wiebe (2000): probability that a sentence is subjective, given that there is at least 1 adjective, is 55.8%.
- ▶ Classifying all sentences as subjective provides us with a much lower precision.
- ▶ Conclusion:
 - ▶ Adjectives, in general, are a good indicator of subjective language.
 - ▶ **But not all adjectives are subjective.**

Why Learning Subjective Adjectives?

- ▶ Try to determine linguistic properties of subjective adjectives.
- ▶ Predictive properties should beat baseline: 55.8%.
- ▶ Research should get us a better understanding of subjective language.
- ▶ In other words: what linguistic properties correlate with subjective language?

What is the point of such an abstract research question?

- ▶ This is **basic research**.
- ▶ The output of such research should *improve scientific theories for improved understanding or prediction of natural or other phenomena* (National Science Foundation).
- ▶ Unlike *applied research*, this type of research does not *immediately* lead to the development of new technology.
- ▶ Basic research is required when researching on an emerging discipline.

What is the point of such an abstract research question?

- ▶ This is **basic research**.
- ▶ The output of such research should *improve scientific theories for improved understanding or prediction of natural or other phenomena* (National Science Foundation).
- ▶ Unlike *applied research*, this type of research does not *immediately* lead to the development of new technology.
- ▶ Basic research is required when researching on an emerging discipline.
- ▶ *When this research was done (~late 90s, early 00s), sentiment analysis was a new research discipline.*

Property to Examine: Polarity

- ▶ Examine whether words that are inherently positive (*beautiful, calm* etc.) or negative (*ugly, noisy* etc.) correlate with subjectivity.
- ▶ Hatzivassiloglou & McKeown (1997) proposed a graph-based approach to extract positive and negative words.
- ▶ Motivation:
 - ▶ Polarity is tightly connected to evaluative judgments.
 - ▶ Evaluative judgements is a subset of the subjectivity (*according to the authors*).

Property to Examine: Gradability

- ▶ Enables a word to participate in comparative constructs:
 - ▶ *small, smaller, smallest, smaller than, as small as*
- ▶ Gradable adjectives express varying degrees of strength relative to a norm (*small planet* is larger than a *large house*).
- ▶ Is it a good predictor of subjectivity?

Property to Examine: Gradability

- ▶ How can we automatically detect gradable adjectives (Hazivassiloglou & Wiebe, 2000):
 - ▶ being modified by specific adverbs and noun phrases:
 - ▶ *a little* <gradable_adj>
 - ▶ *exceedingly* <gradable_adj>
 - ▶ *very* <gradable_adj>
 - ▶ adjective that is inflected:
 - ▶ *larger, largest*
 - ▶ *uglier, ugliest*
- ▶ Run these patterns on a large (unlabelled) corpus and collect gradable adjectives.

Why choosing *polarity* and *gradability* as properties to inspect?

- ▶ Both are linguistic concepts.
- ▶ For both properties, there already exist lexical resources to be used as data for this research.
- ▶ Previous research successfully established *automatic* extraction methods.
- ▶ Would not be very convincing to propose a detection method based on properties which cannot be automatically detected.

Evaluation Set Up

- ▶ Have a section of the Wall Street Journal Treebank Corpus *manually* annotated.
- ▶ Size: ~1000 sentences
- ▶ Sentence-level annotation: subjective/objective → for task evaluation.
- ▶ Classify a sentence as subjective if a particular adjective is contained (e.g. gradable adjective).
- ▶ Word-level annotation: each subjective word on the subjective sentences was additionally marked → for baseline.

Evaluation Set Up (Baseline)

- ▶ Divide annotated sentences into training and test data.
- ▶ Create seeds:
 - ▶ Get all subjective adjectives from training set.
- ▶ Test the output on the test data:
 - ▶ A sentence is predicted as subjective, if at least one subjective adjective (according to training data) is observed.
 - ▶ On the set of adjectives predicted: compute precision of being subjective.

Some Further Details

- ▶ Seeds are not evaluated as such: adjectives that are *distributionally similar* are also included.
- ▶ Experiments done on 10-fold crossvalidation:
 - ▶ Unusual partitioning:
 - ▶ 1/10 training
 - ▶ 9/10 test

Evaluation

Feature	Precision
Baseline I: 1 arbitrary adj.	55.8
Baseline II: Seeds	63.3
Polarity	64.3
Seeds & Polarity	67.6
Gradability	68.2
Seeds & Gradability	79.6

Conclusion

- ▶ Subjective adjectives can be learned from annotated data.
- ▶ Established two predictive criteria:
 - ▶ polarity
 - ▶ gradability

Outline

- ▶ Learning Subjective Adjectives
- ▶ Learning Subjective Nouns using Extraction Pattern Bootstrapping

“Learning Subjective Nouns using Extraction Pattern Bootstrapping”

**By Ellen Riloff, Janyce Wiebe, Theresa Wilson,
In *Proceedings of CoNLL*, 2003.**

Objective

- ▶ Learn a large set of subjective nouns.
- ▶ Employ pattern bootstrapping.
- ▶ Generic lexicon extraction algorithm.
- ▶ Evaluate extracted subjective nouns in detection of subjective sentences.

Data

- ▶ All data are English-language versions of foreign news documents.
- ▶ Annotation is at sentence level:
 - ▶ *subjective vs. objective*
- ▶ 109 documents for training and testing approach.
 - ▶ 2197 sentences
- ▶ 33 documents as a development set (*tuning corpus*).
 - ▶ 698 sentences

Interannotation Agreement

- ▶ 12 documents (178 sentences) were annotated in parallel.
- ▶ Kappa value of 0.71 → substantial agreement.

What is Bootstrapping?

In computing, **bootstrapping** refers to a process where a simple system activates another more complicated system that serves the same purpose. It is a solution to the *chicken-and-egg problem* of starting a certain system without the system already functioning.

What is Bootstrapping?

In computing, **bootstrapping** refers to a process where a **simple system** activates another more complicated system that serves the same purpose. It is a solution to the *chicken-and-egg problem* of starting a certain system without the system already functioning.

What is Bootstrapping?

In computing, **bootstrapping** refers to a process where a **simple system (=seeds for subjective nouns)** activates another more complicated system that serves the same purpose. It is a solution to the *chicken-and-egg problem* of starting a certain system without the system already functioning.

What is Bootstrapping?

In computing, **bootstrapping** refers to a process where a simple system activates another more complicated system that serves the same purpose. It is a solution to the *chicken-and-egg problem* of starting a certain system without the system already functioning.

What is Bootstrapping?

In computing, **bootstrapping** refers to a process where a simple system activates another more **complicated system** that serves the same purpose. It is a solution to the *chicken-and-egg problem* of starting a certain system without the system already functioning.

What is Bootstrapping?

In computing, **bootstrapping** refers to a process where a simple system activates another more **complicated system (=pattern-based system to extract subjective nouns)** that serves the same purpose. It is a solution to the *chicken-and-egg problem* of starting a certain system without the system already functioning.

Bootstrapping: Origins of the Term

Bootstrapping alludes to a German legend about a **Baron Muenchhausen**, who was able to lift himself out of a swamp by pulling himself up by his own hair (*see picture on the right*).



Bootstrapping: Origins of the Term

In later versions he was using his own **bootstraps** to pull himself out of the sea.



The Basilisk Bootstrapping Algorithm

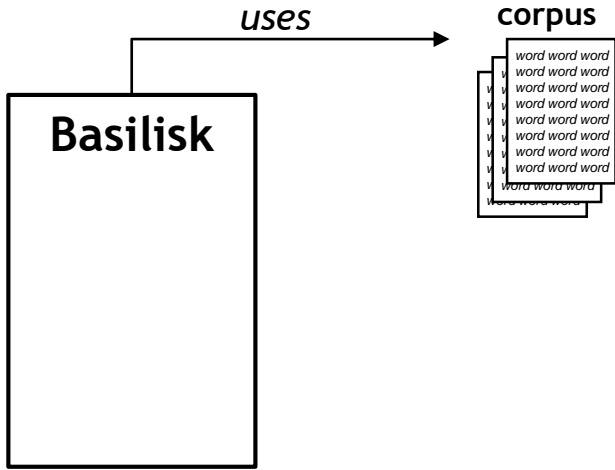


▶ Input:

- ▶ A set of seed words (i.e. subjective nouns).
- ▶ A large unannotated corpus (950 different news texts from the same collection as the annotated data).

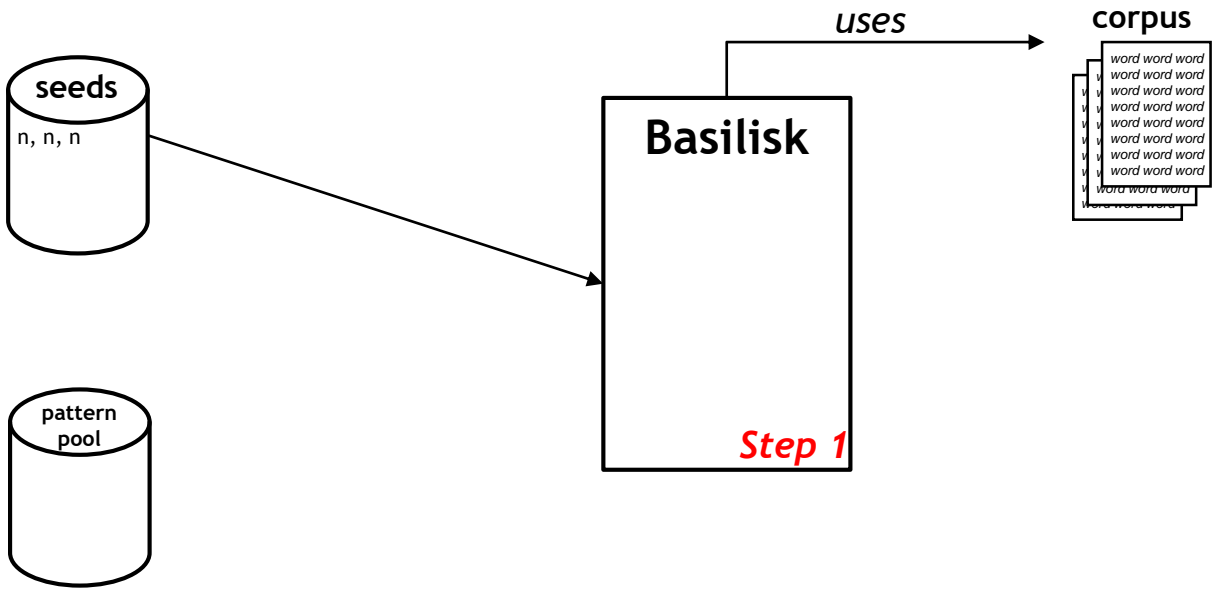
▶ Output:

- ▶ A set of (potentially) subjective nouns.
- ▶ A set of extraction patterns by which these nouns have been found:
 - ▶ Example: *voiced* <*dobj*> (as in *voiced concern*)
 - ▶ where *dobj* is a direct object representing a subjective noun

The Basilisk Bootstrapping Algorithm

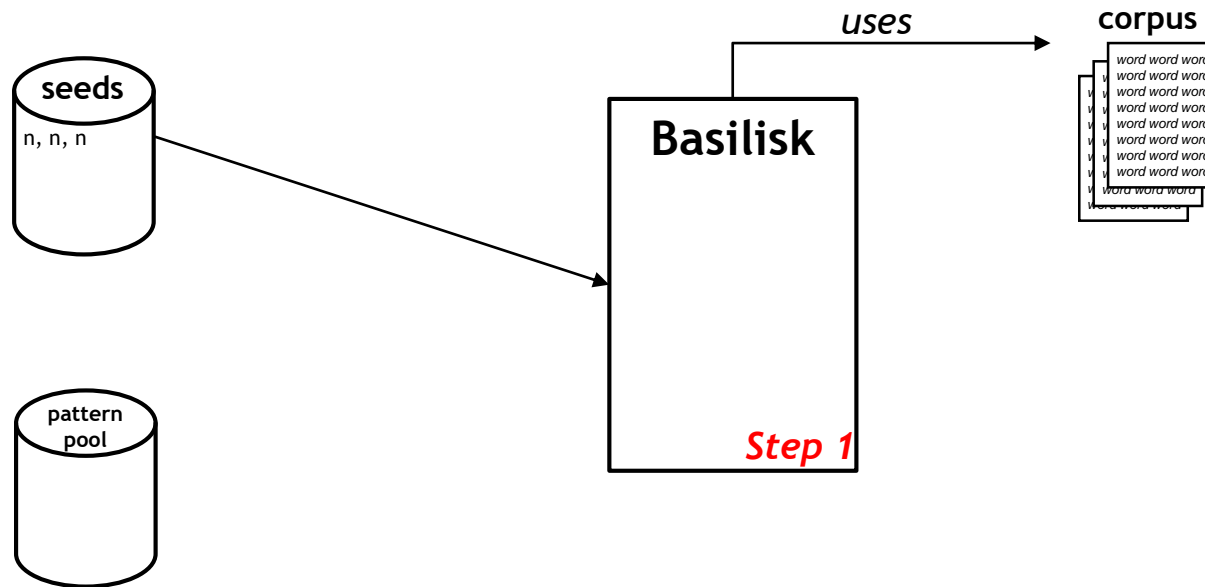


The Basilisk Bootstrapping Algorithm



iteration 0


The Basilisk Bootstrapping Algorithm



- computes patterns with the help of seeds
- patterns are contexts that correlate with the seeds

iteration 0

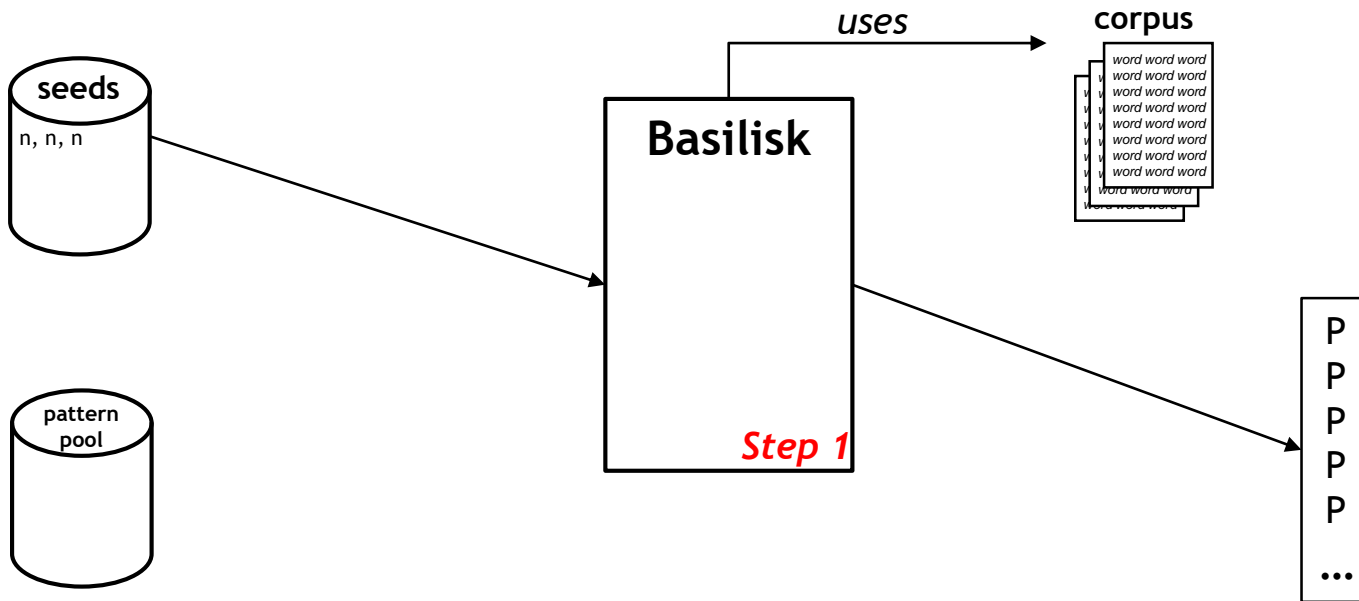
The Basilisk Bootstrapping Algorithm



- computes patterns with the help of seeds
- patterns are contexts that correlate with the seeds

iteration 0

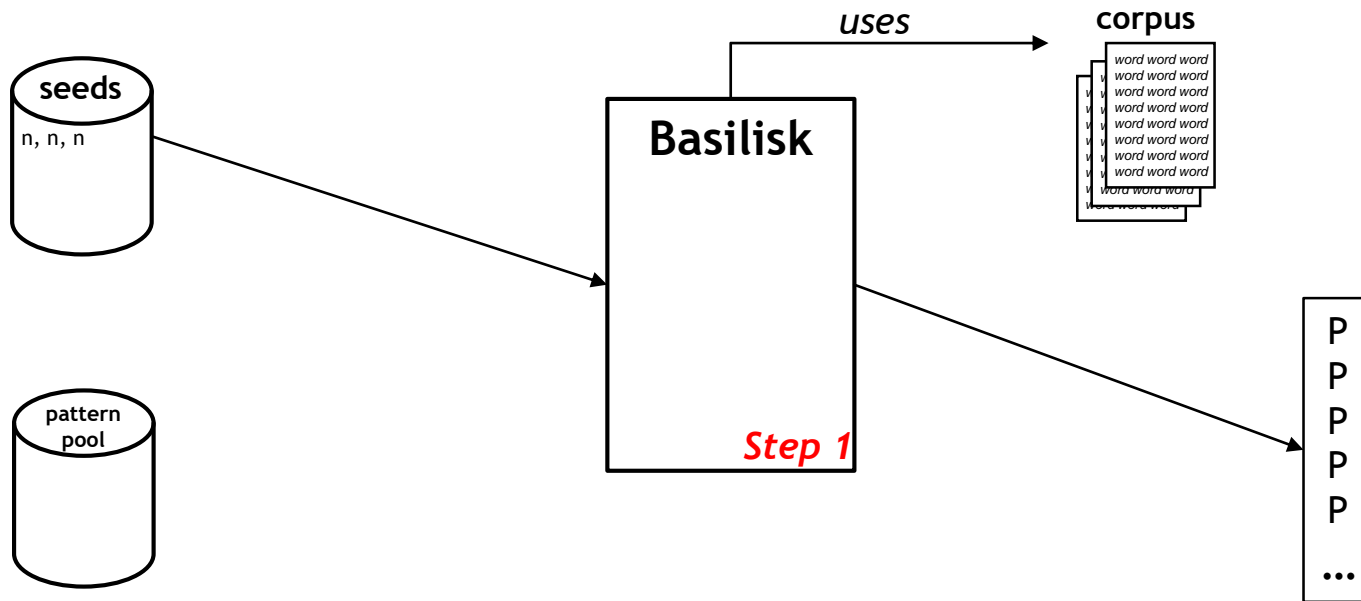
The Basilisk Bootstrapping Algorithm



- computes patterns with the help of seeds
- patterns are contexts that correlate with the seeds
- patterns are ranked according to correlation

iteration 0

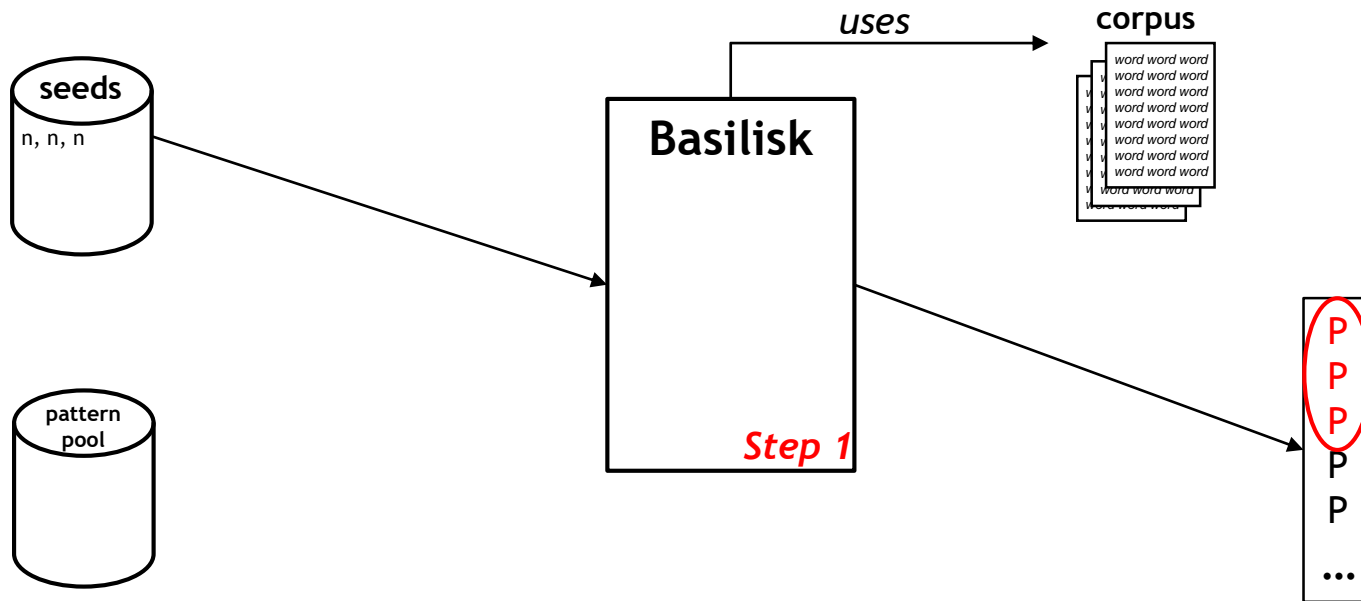
The Basilisk Bootstrapping Algorithm



- computes patterns with the help of seeds
- patterns are contexts that correlate with the seeds
- patterns are ranked according to correlation
- best patterns are added to pattern pool

iteration 0

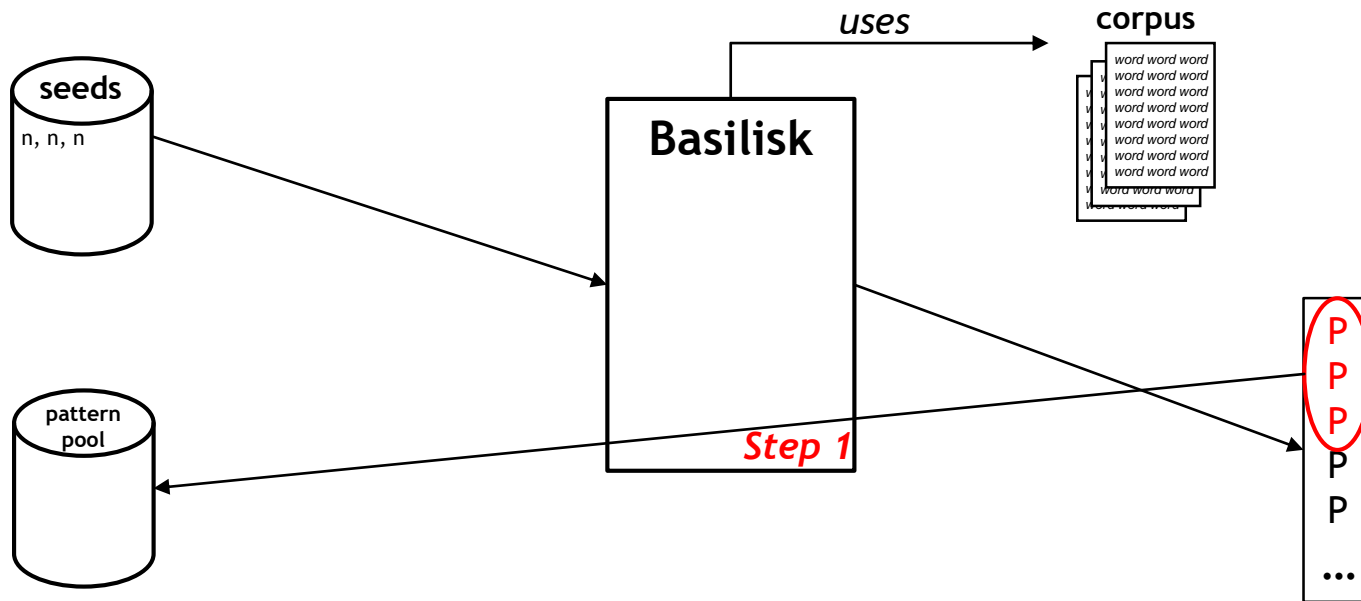
The Basilisk Bootstrapping Algorithm



- computes patterns with the help of seeds
- patterns are contexts that correlate with the seeds
- patterns are ranked according to correlation
- best patterns are added to pattern pool

iteration 0

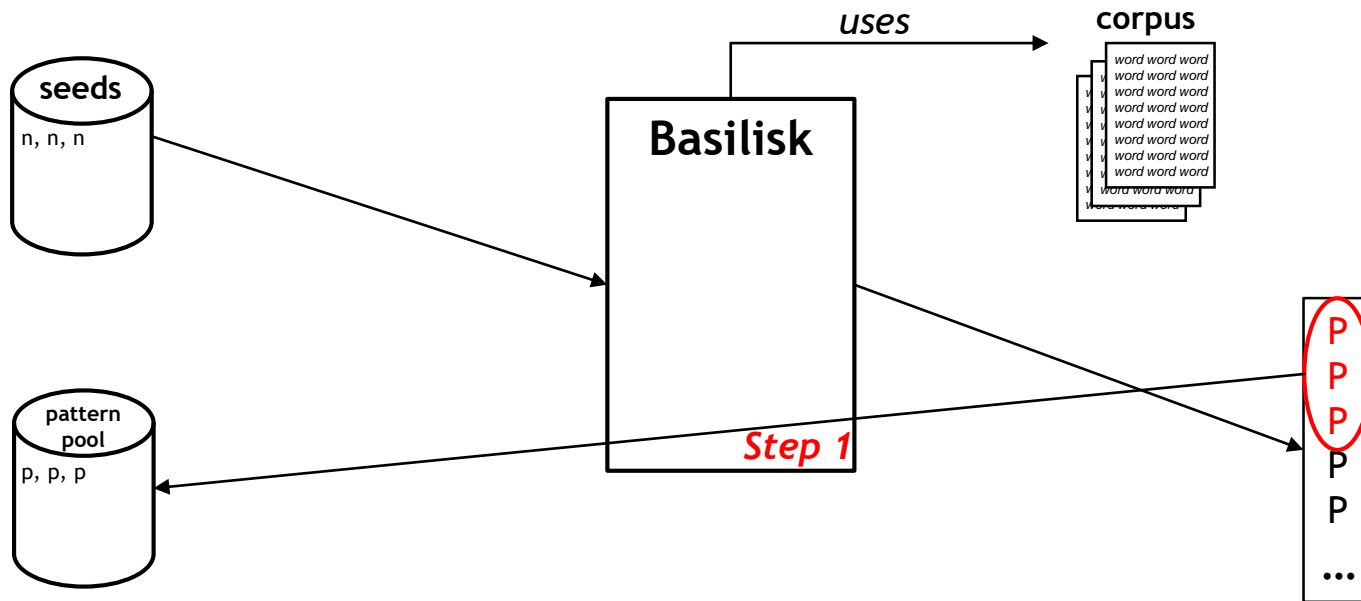
The Basilisk Bootstrapping Algorithm



- computes patterns with the help of seeds
- patterns are contexts that correlate with the seeds
- patterns are ranked according to correlation
- best patterns are added to pattern pool

iteration 0

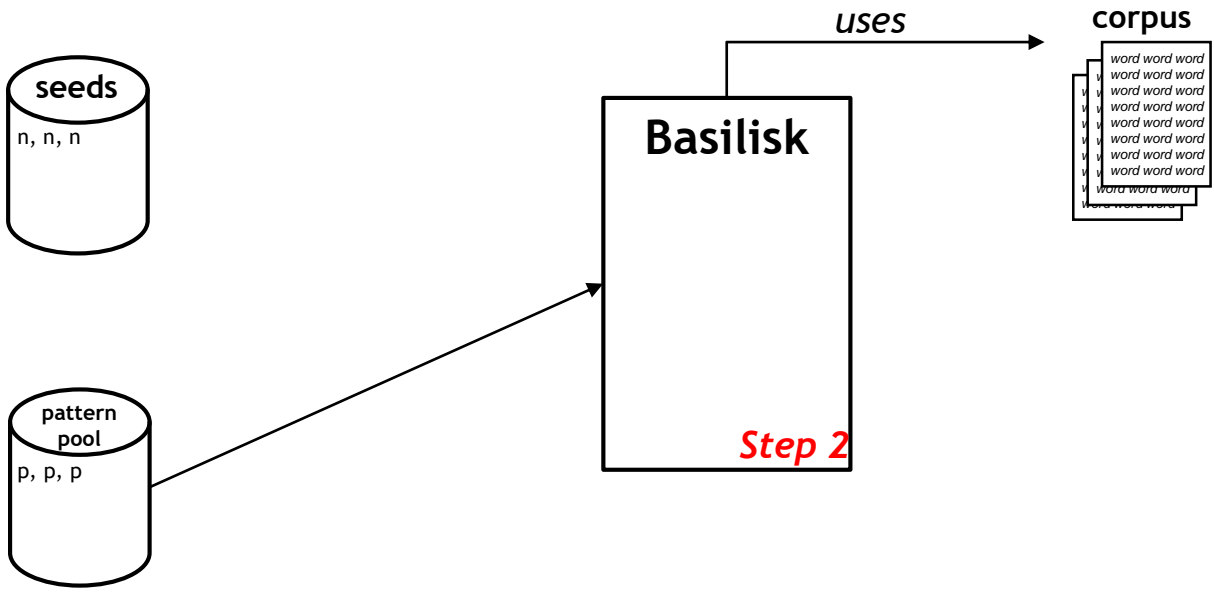
The Basilisk Bootstrapping Algorithm



- computes patterns with the help of seeds
- patterns are contexts that correlate with the seeds
- patterns are ranked according to correlation
- best patterns are added to pattern pool


iteration 0

The Basilisk Bootstrapping Algorithm



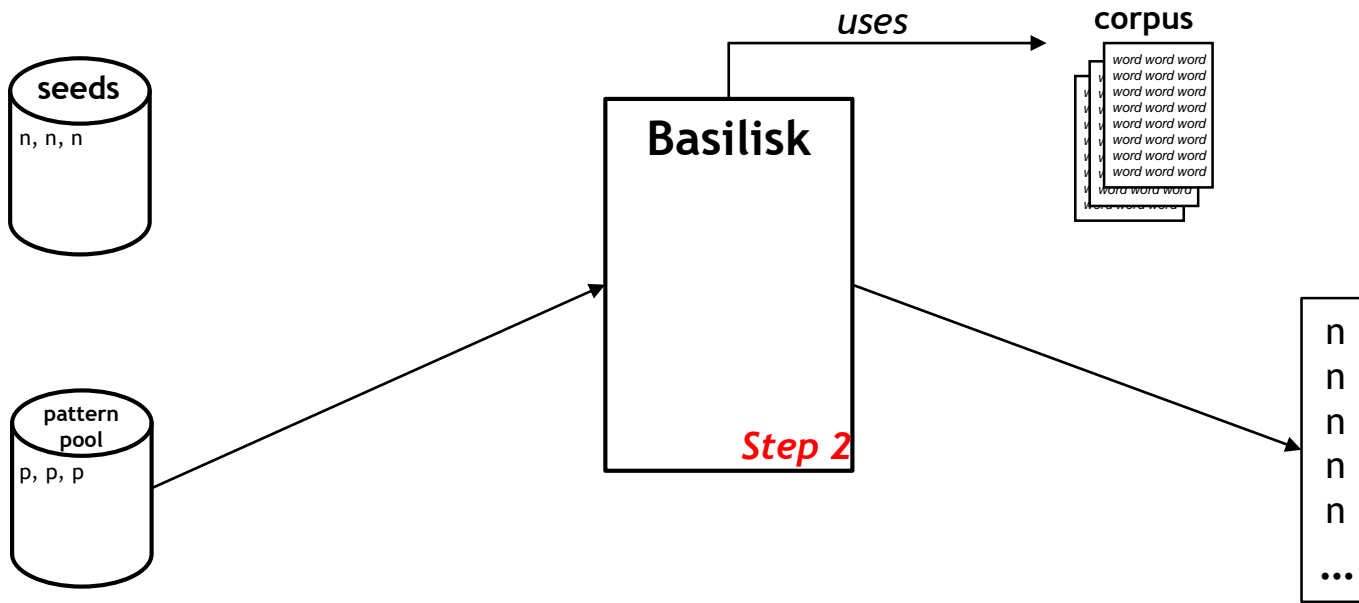
iteration 0

The Basilisk Bootstrapping Algorithm



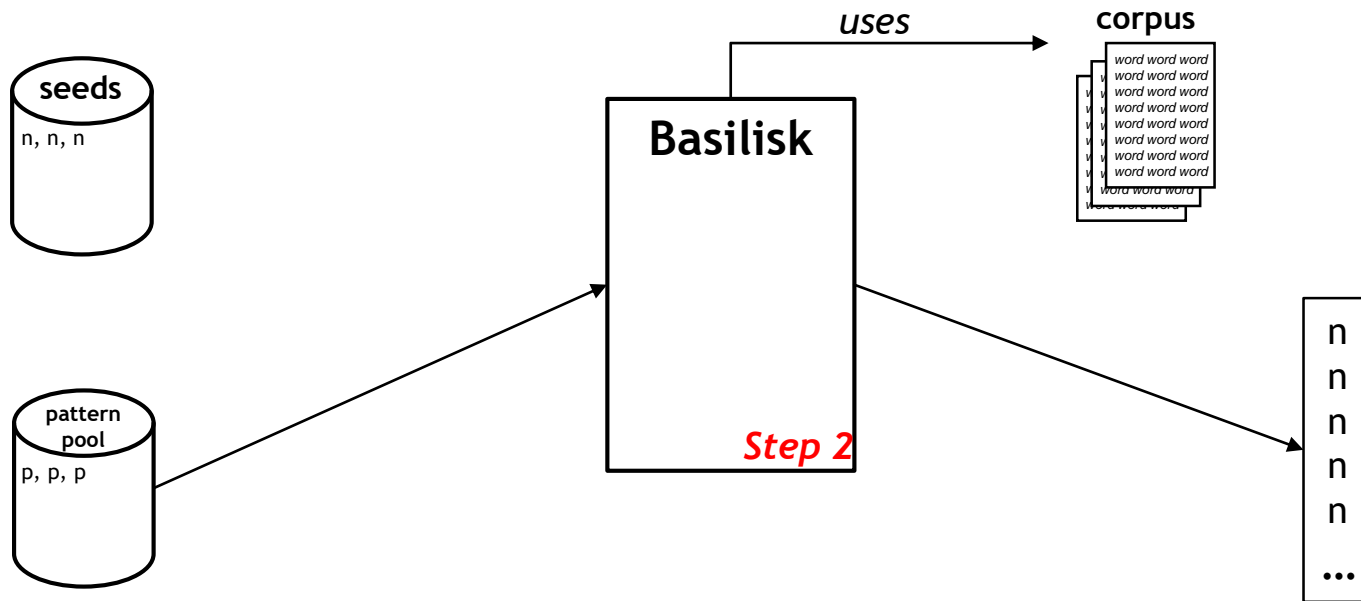
- extract nouns that can be found with patterns

The Basilisk Bootstrapping Algorithm



- extract nouns that can be found with patterns


The Basilisk Bootstrapping Algorithm



- extract nouns that can be found with patterns
- nouns are scored according to how well they correlate with patterns

iteration 0

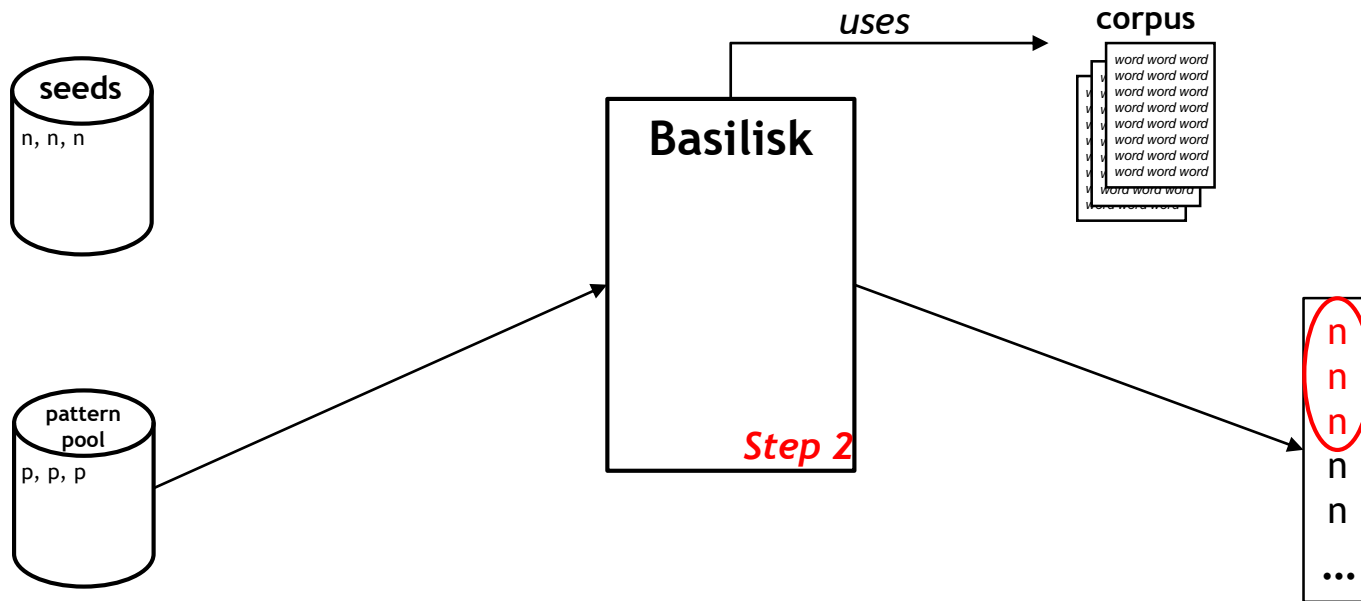
The Basilisk Bootstrapping Algorithm



- extract nouns that can be found with patterns
- nouns are scored according to how well they correlate with patterns
- add top m nouns to seeds

iteration 0

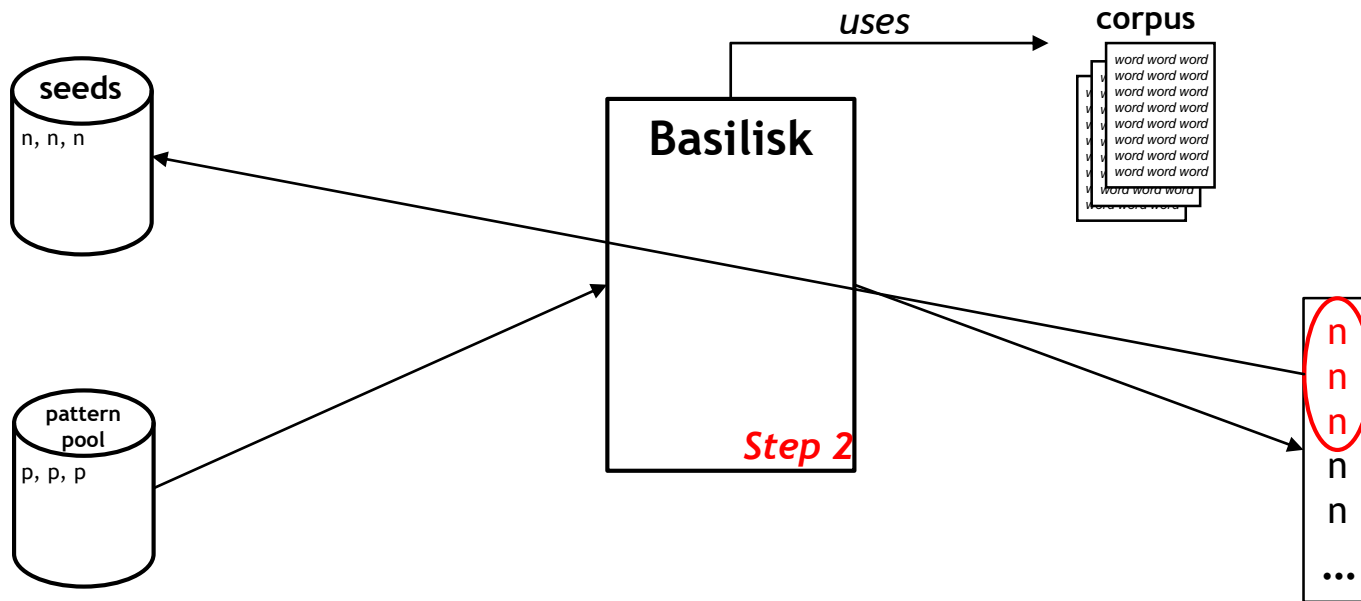
The Basilisk Bootstrapping Algorithm



- extract nouns that can be found with patterns
- nouns are scored according to how well they correlate with patterns
- add top m nouns to seeds

iteration 0

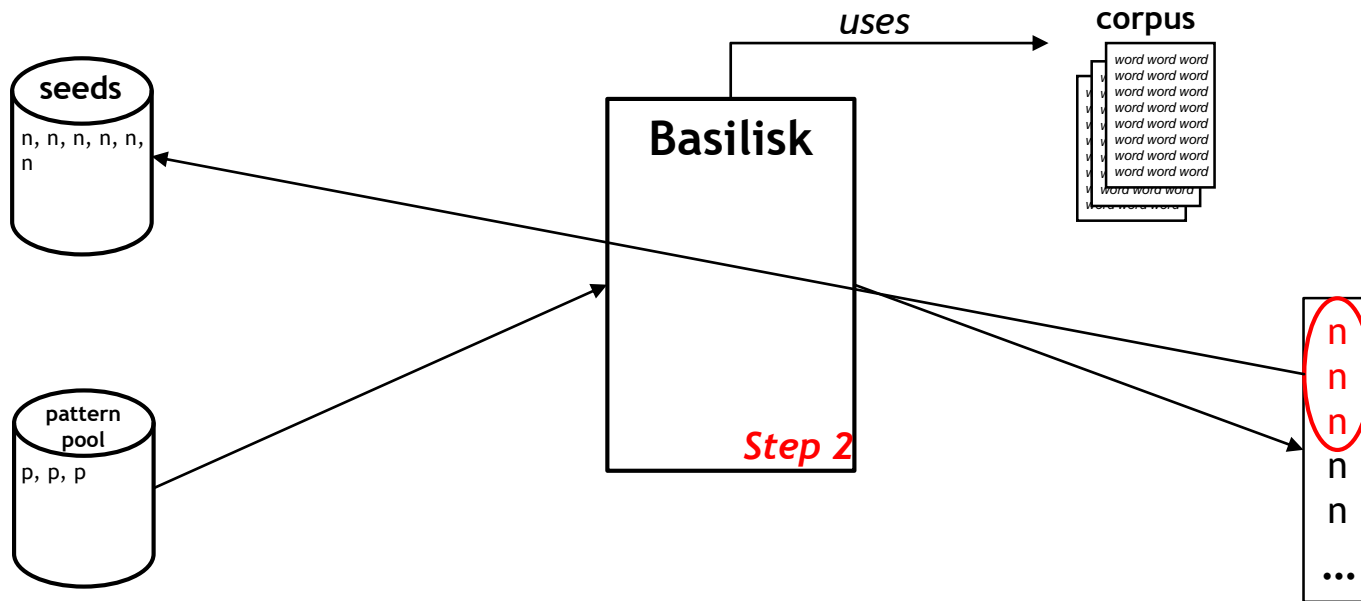
The Basilisk Bootstrapping Algorithm



- extract nouns that can be found with patterns
- nouns are scored according to how well they correlate with patterns
- add top m nouns to seeds

iteration 0

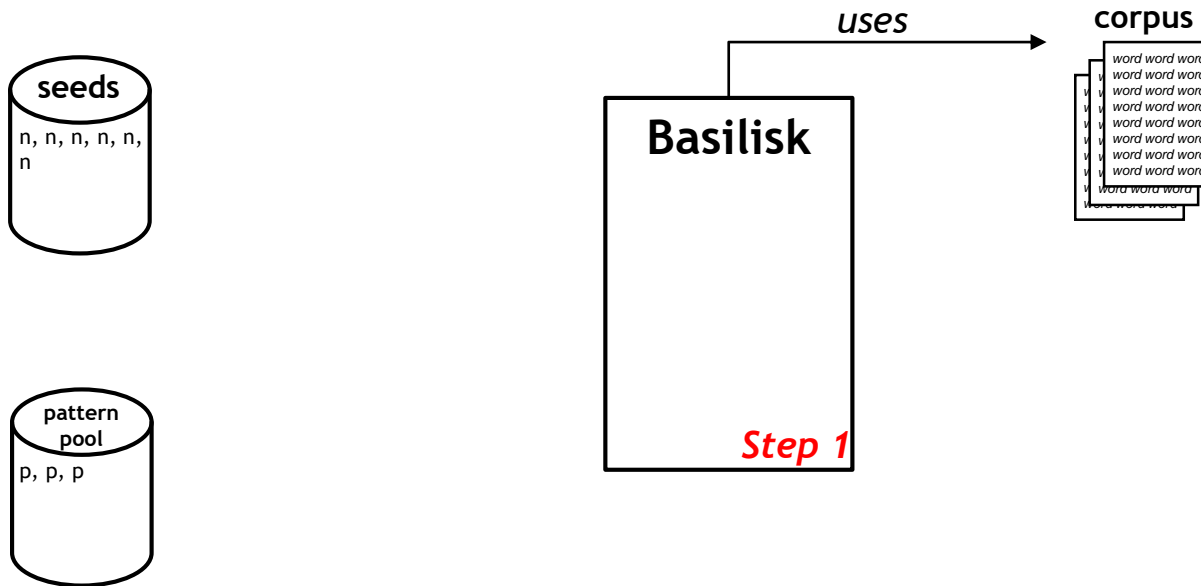
The Basilisk Bootstrapping Algorithm



- extract nouns that can be found with patterns
- nouns are scored according to how well they correlate with patterns
- add top m nouns to seeds


iteration 0

The Basilisk Bootstrapping Algorithm



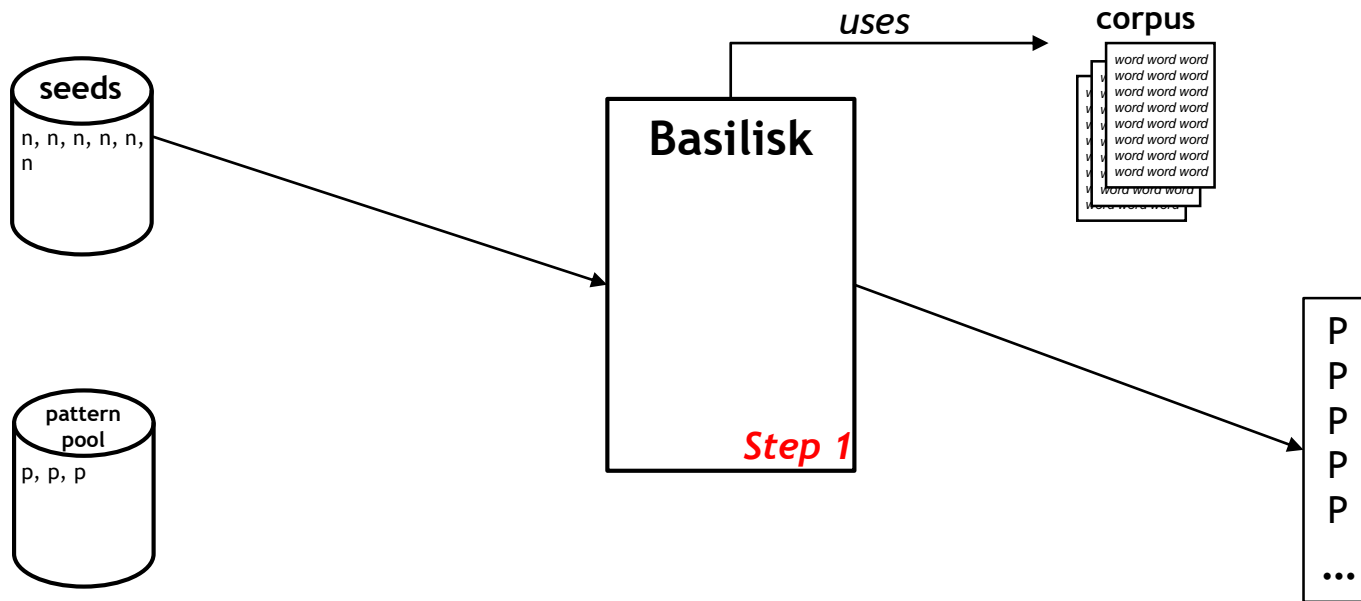
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



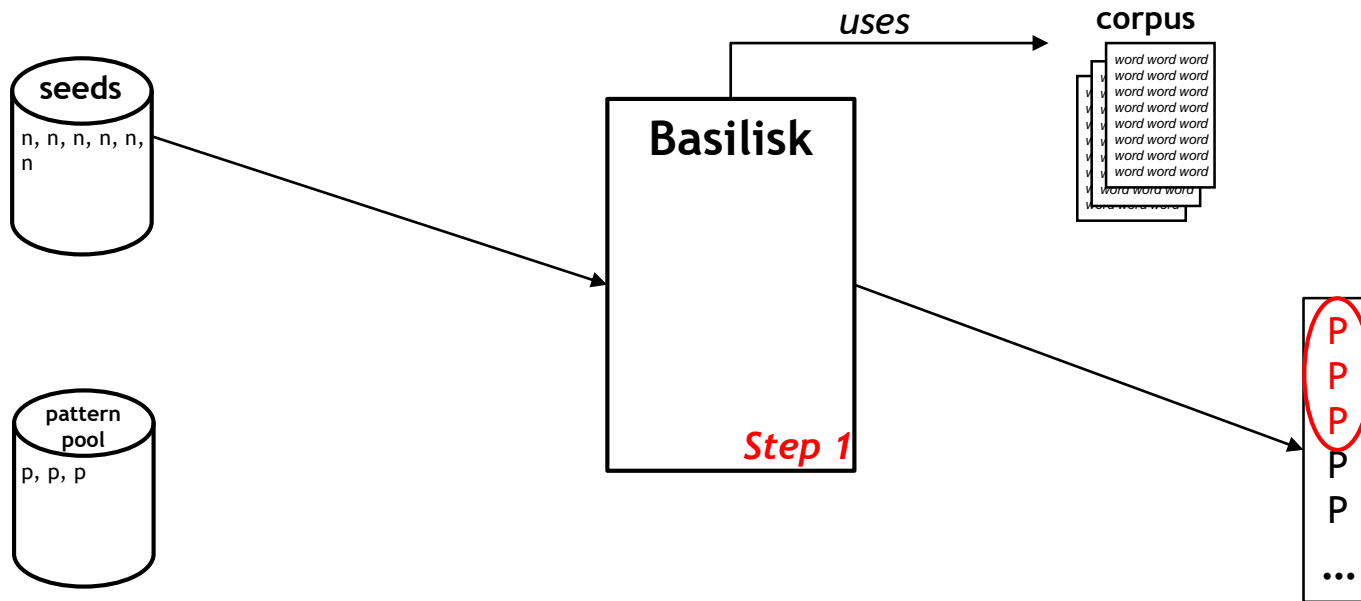
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



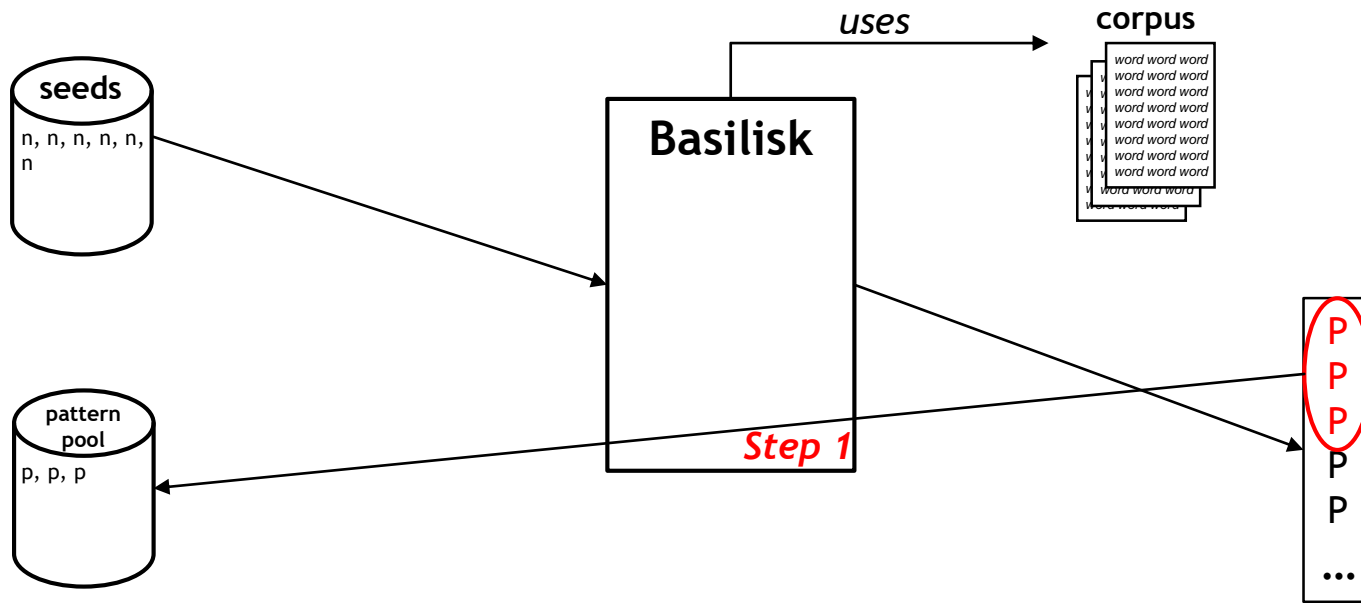
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



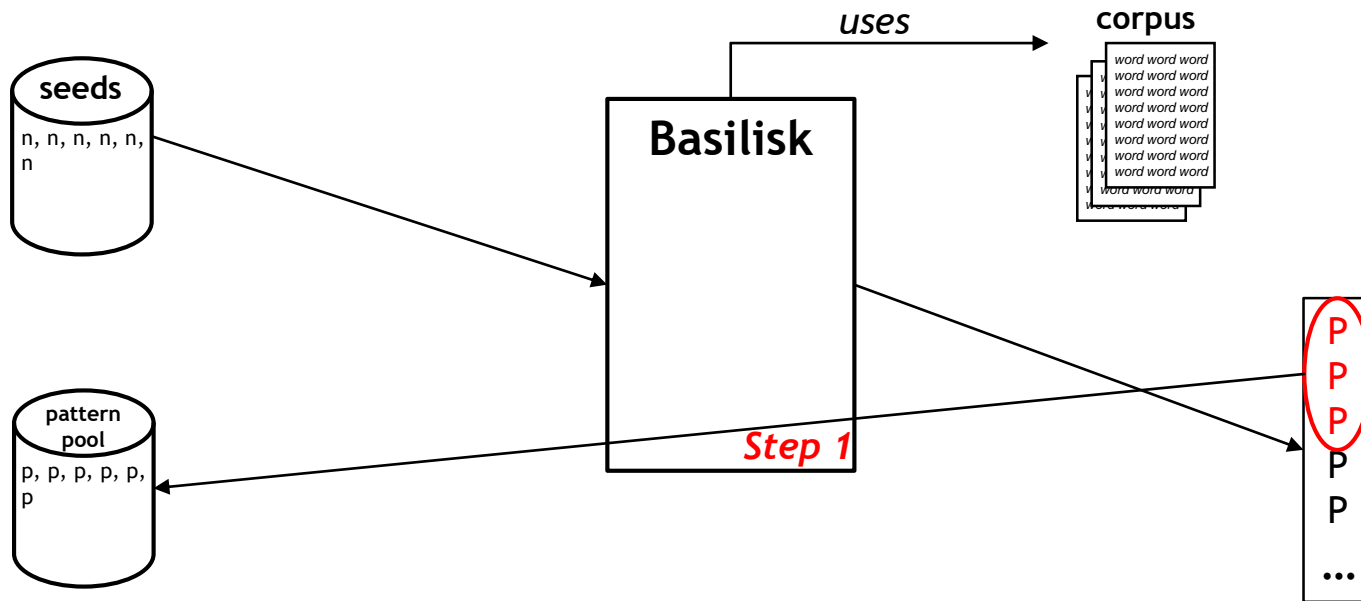
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



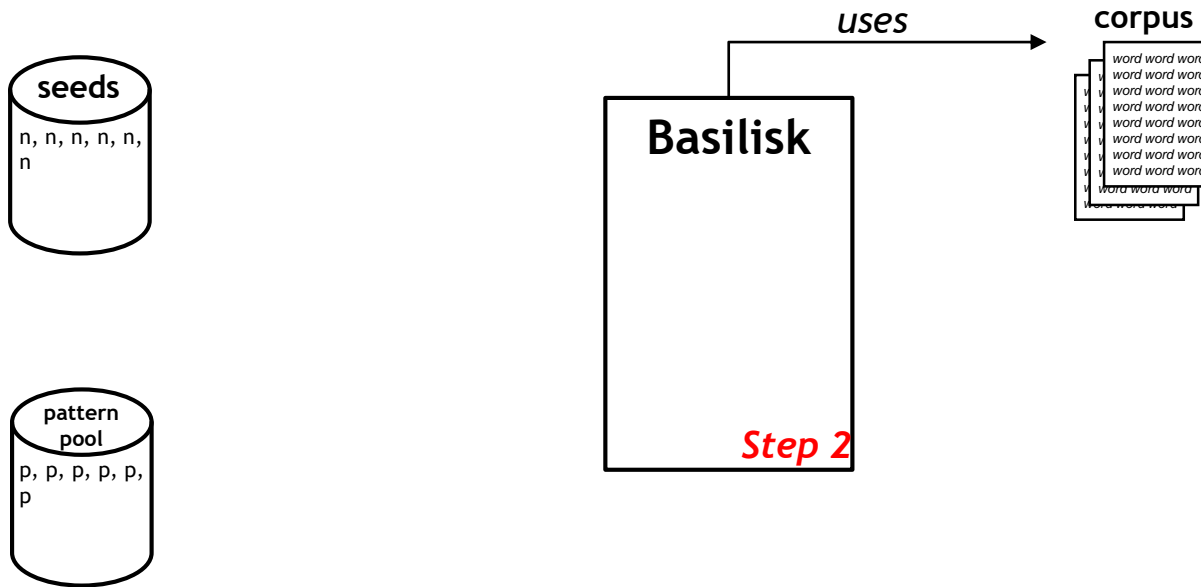
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



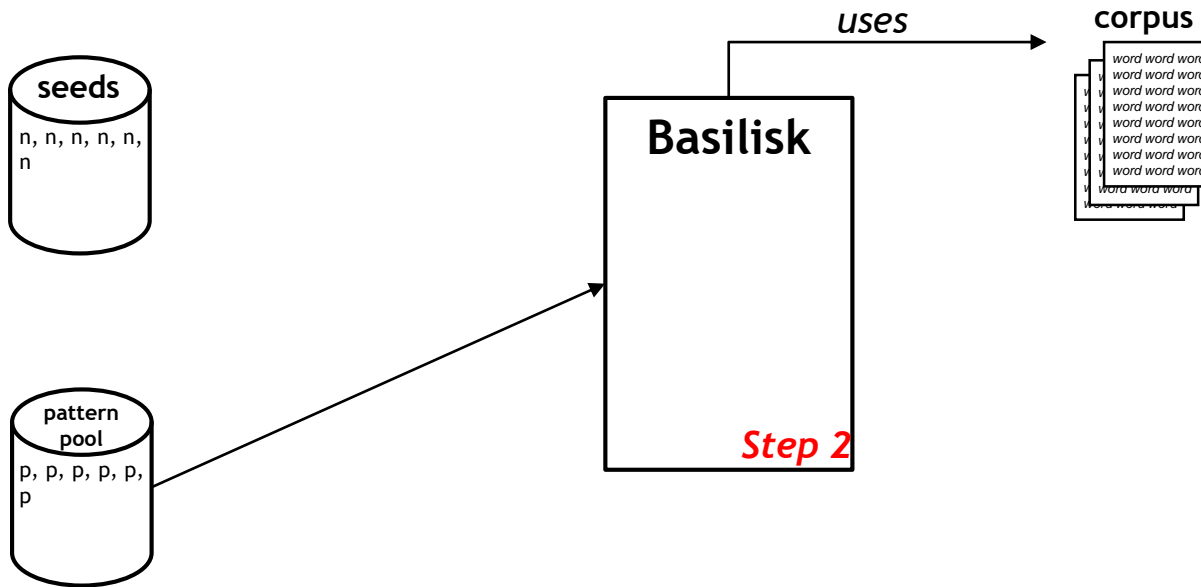
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



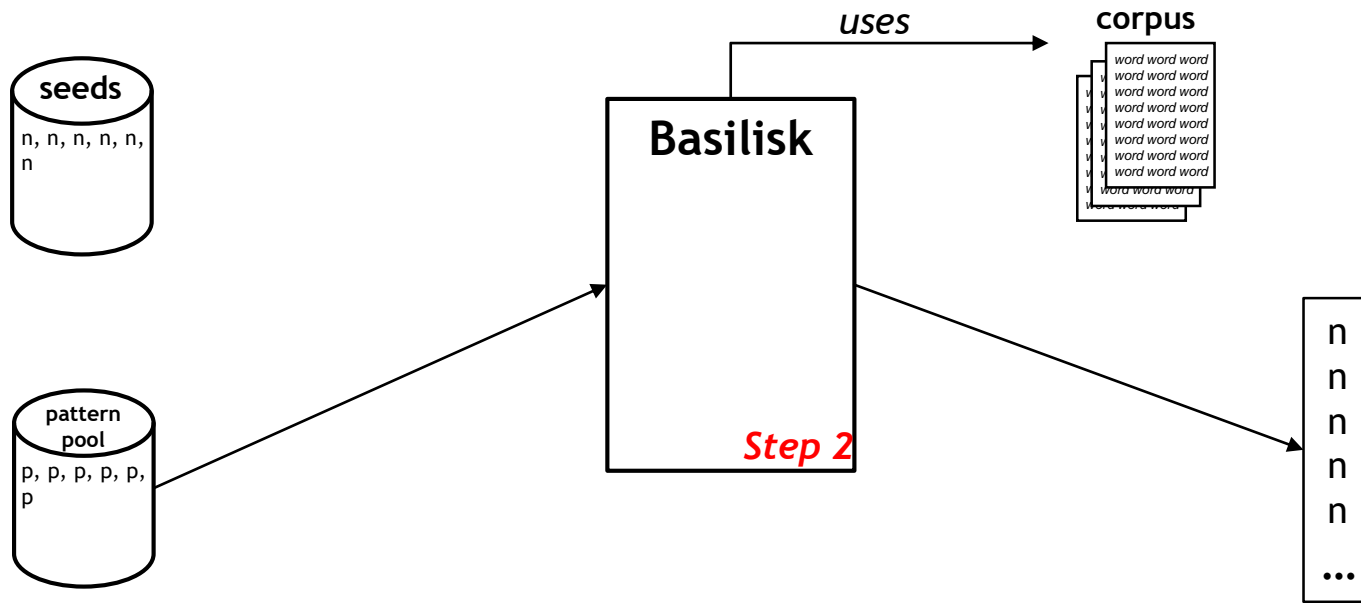
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



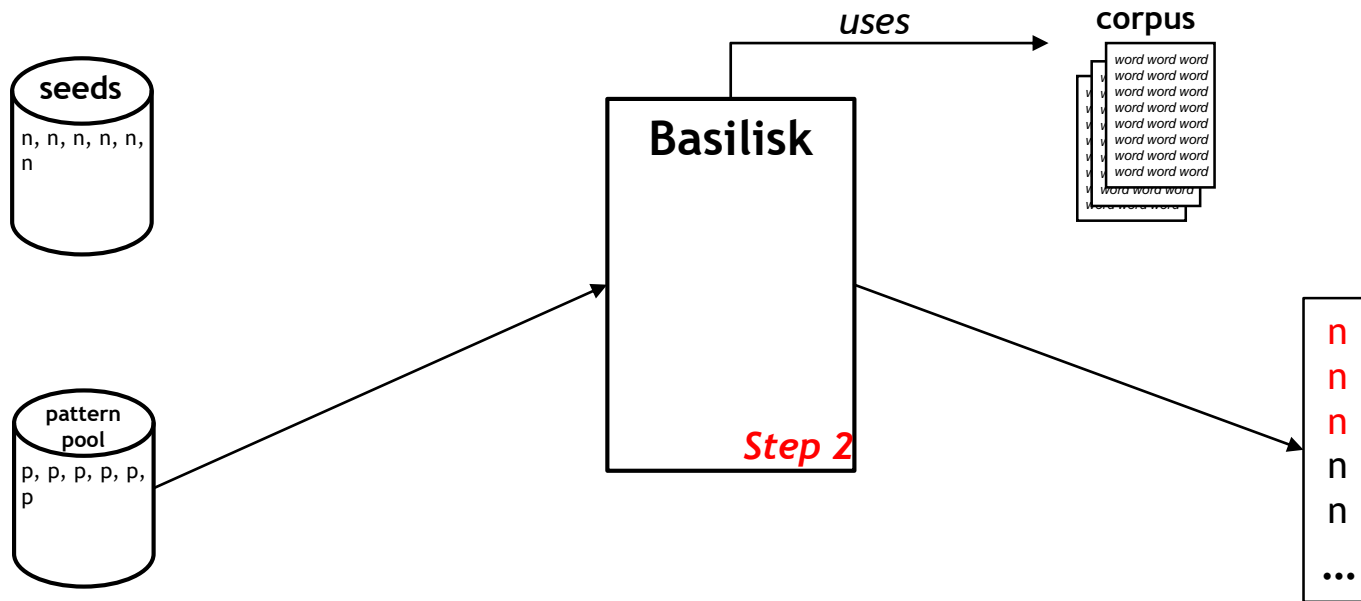
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



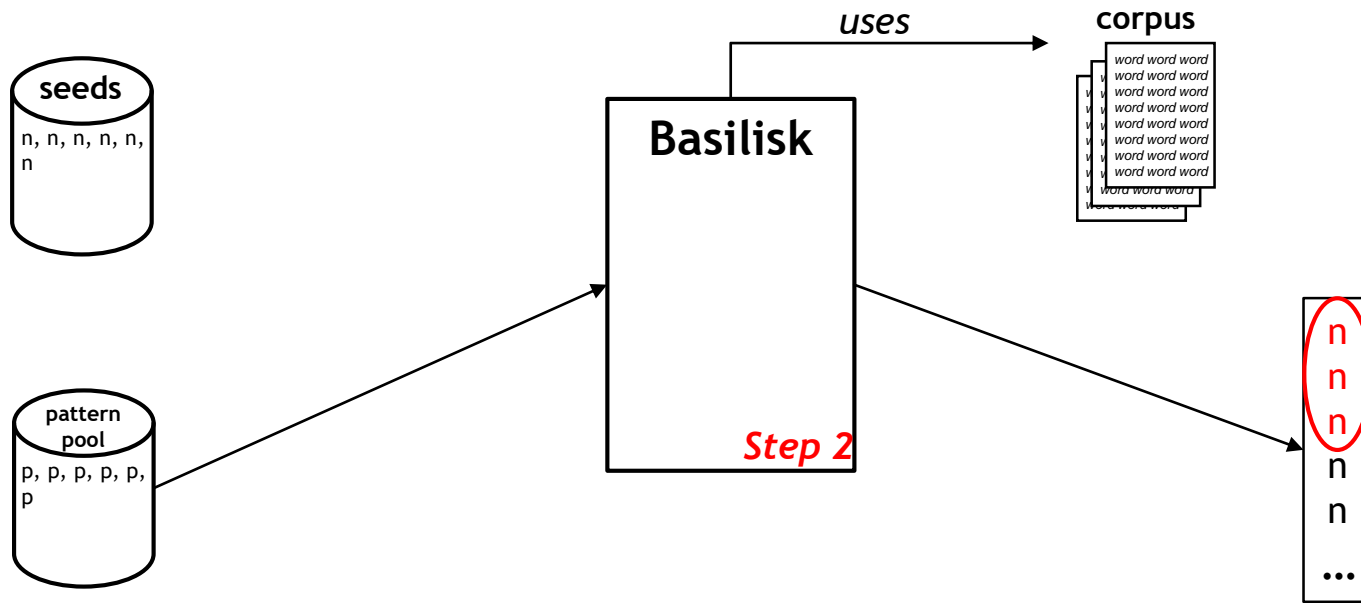
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



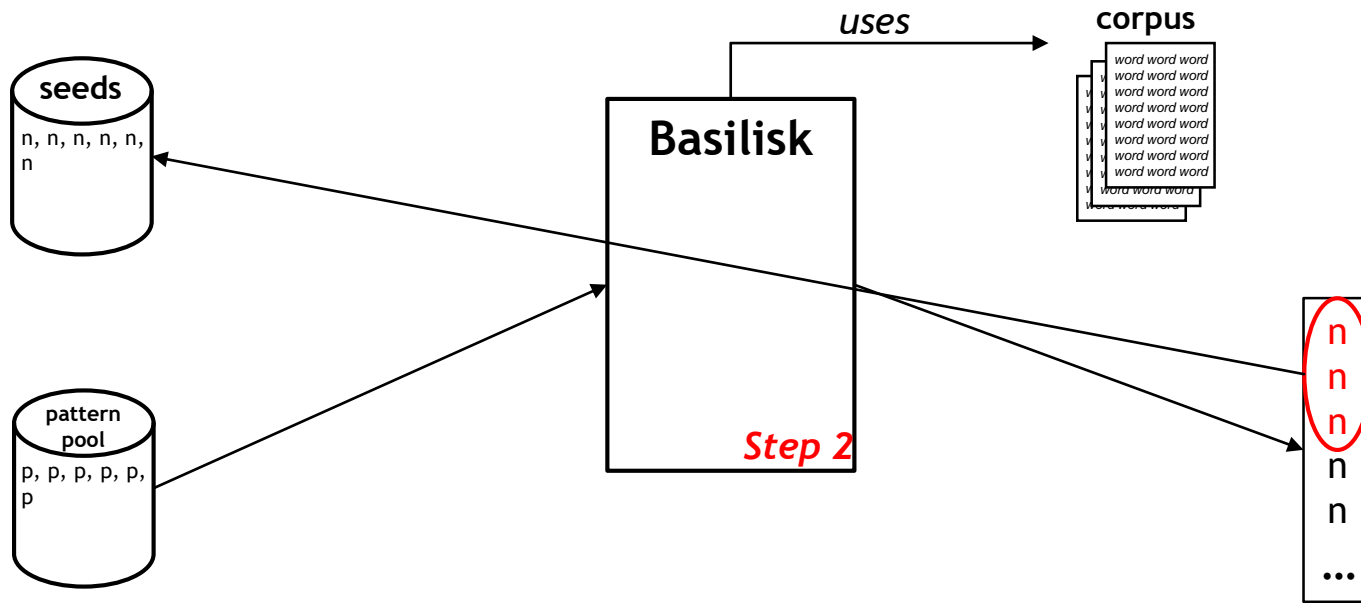
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



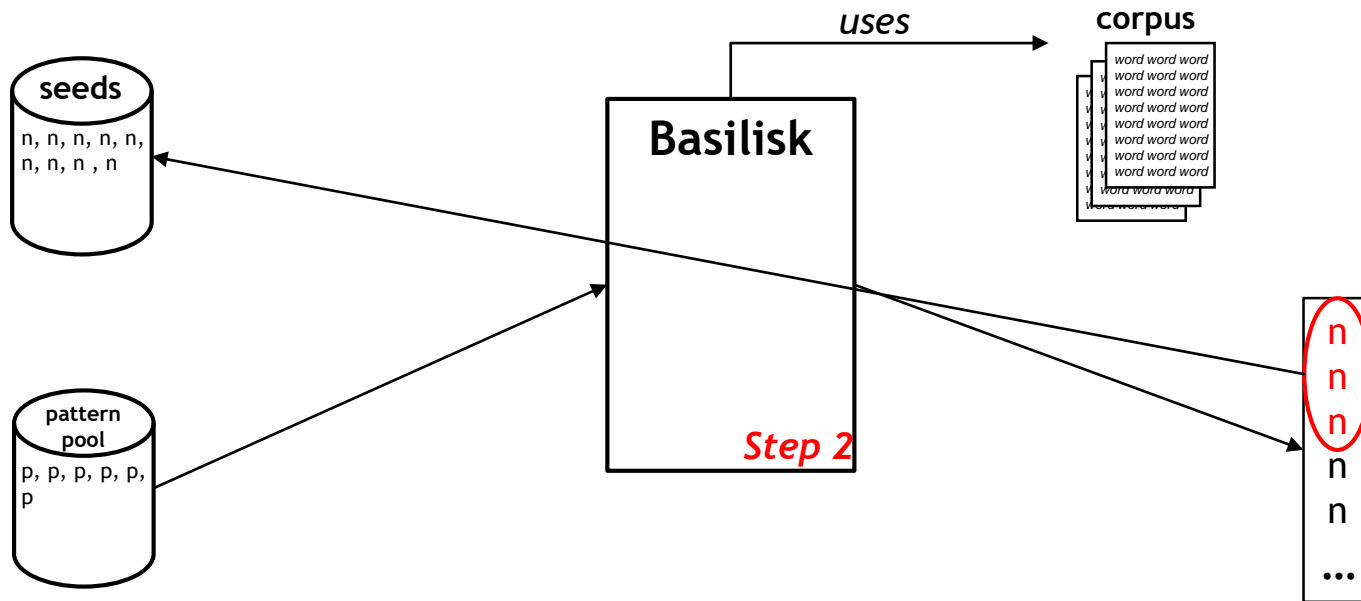
- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm




- We repeat steps 1 and 2 for an empirically determined amount of times!

The Basilisk Bootstrapping Algorithm



- We repeat steps 1 and 2 for an empirically determined amount of times!

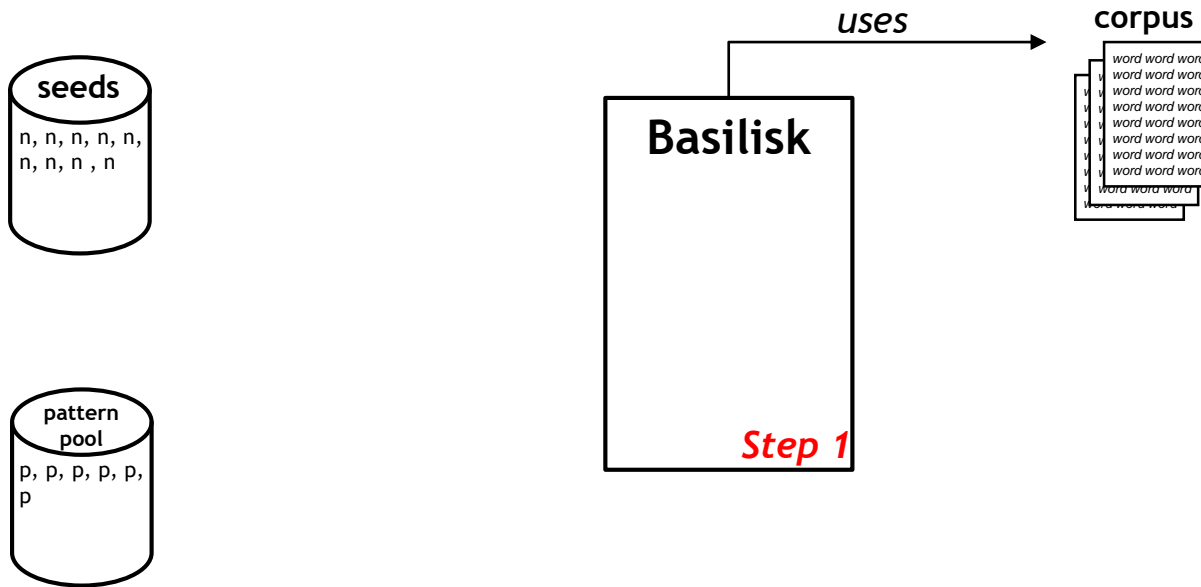
The Basilisk Bootstrapping Algorithm



- We repeat steps 1 and 2 for an empirically determined amount of times!

iteration 2

The Basilisk Bootstrapping Algorithm



- We repeat steps 1 and 2 for an empirically determined amount of times!
- In the present setting: 400 iterations adding 5 words at each iteration.

iteration 2

Seed Selection

- ▶ 850 nouns positively correlated with subjective sentences.
- ▶ Sentences are not from the present gold standard!
- ▶ Sort words by frequency.
- ▶ Manually select a set of 20 high-frequency words considered to be *strongly subjective*.

Illustration of Seeds

cowardice	embarrassment	hatred	outrage
crap	fool	hell	slander
delight	gloom	hypocrisy	sigh
disdain	grievance	love	twit
dismay	happiness	nonsense	virtue

Illustration of Learned Patterns

Extraction Patterns	Examples of Extracted Nouns
expressed <dobj>	condolences, hope, grief, views, worries, recognition
indicative of <np>	compromise, desire, thinking
inject <dobj>	vitality, hatred
reaffirmed <dobj>	resolve, position, commitment
voiced <dobj>	outrage, support, skepticism, disagreement, opposition, concerns, gratitude, indignation
show of <np>	support, strength, goodwill, solidarity, feeling
<subject> was shared	anxiety, view, niceties, feeling

Manual Review of the Output

- ▶ Output of Bootstrapping: ~2000 nouns.
- ▶ Each word is categorized as either:
 - ▶ **Strong Subjective:** strong subjective connotation/unambiguous (e.g. *bully*, *barbarian*)
 - ▶ **Weak Subjective:** weak subjective connotation/ambiguous (e.g. *aberration*, *plague*)
 - ▶ **Objective**
- ▶ **After manual review: 825 subjective nouns**
 - ▶ Strong Subjective: 372 nouns
 - ▶ Weak Subjective: 453 nouns

Examples of Learned Subjective Nouns

Strong Subjective		Weak Subjective	
tyranny	scum	aberration	plague
smokescreen	bully	allusion	risk
apologist	devil	apprehensions	drama
barbarian	liar	beneficiary	trick
belligerence	pariah	resistant	promise
condemnation	venom	credence	intrigue
sanctimonious	diatribe	distortion	unity
exaggeration	mockery	eyebrows	failures
repudiation	anguish	inclination	tolerance
insinuation	fallacies	liability	persistent
antagonism	evil	assault	trust
atrocities	genius	benefit	success
denunciation	goodwill	blood	spirit
exploitation	injustice	controversy	slump
humiliation	innuendo	likelihood	sincerity
ill-treatment	revenge	peaceful	eternity
sympathy	rogue	pressure	rejection

What's the point of an automatic extraction if it requires manual reviewing?

- ▶ Semi-automatic process is faster than manual annotation:
 - ▶ Semi-automatic process: more than 40% of the given list is subjective.
 - ▶ Manual annotation: random sample of nouns would have a much lower proportion of subjective nouns.
- ▶ Don't forget, we have an interesting by-product:
 - ▶ The extraction patterns.

Evaluation of the Subjective Nouns: Sentence Classification

- ▶ Do the extracted subjective nouns improve a classifier which detects subjective sentences?
- ▶ Train a supervised classifier (Naive Bayes) using state-of-the-art features (*at that time!*).
- ▶ Task: distinguish between subjective and objective sentences.
- ▶ Train with and without subjective nouns as a feature:
 - ▶ Features:
 - ▶ How many Strong Subjective-nouns are in sentence?
 - ▶ How many Weak Subjective-nouns are in sentence?
- ▶ Do 25-fold cross-validation using the document collection with labeled sentences.

Evaluation of the Subjective Nouns: Sentence Classification

- ▶ Other features:
 - ▶ Bag of words
 - ▶ “WBO” (Wiebe et al., 1999):
 - ▶ Frequency of words correlating with subjective language.
 - ▶ Frequency of words correlating with objective language.
 - ▶ Presence of pronoun/adjective/adverb/cardinal number/modal in sentence.
 - ▶ manual:
 - ▶ Frequency of words according to manually designed lexicons of subjective words.
 - ▶ discourse:
 - ▶ Features that address the density of subjective content in the context (previous/next sentence).
 - ▶ Are there more subjective words in the context than on average in the dataset?

Evaluation of the Subjective Nouns: Sentence Classification

Classifier	Acc	Prec	Rec	F1
<i>Majority Class Classifier</i>	59.0	59.0	100.0	74.2
Bag of Words	73.3	81.7	70.9	75.9
WBO	72.1	76.0	77.4	76.7
WBO+SubjNoun	74.3	78.6	77.8	78.2
WBO+SubjNoun>manual+discourse	76.1	81.3	77.4	79.3

Evaluation of the Subjective Nouns: Sentence Classification

Classifier	Acc	Prec	Rec	F1
<i>Majority Class Classifier</i>	59.0	59.0	100.0	74.2
Bag of Words	73.3	81.7	70.9	75.9
WBO	72.1	76.0	77.4	76.7
WBO+SubjNoun	74.3	78.6	77.8	78.2
WBO+SubjNoun>manual+discourse	76.1	81.3	77.4	79.3

Differences in Accuracies (with the exception of *WBO* and *Bag of Words*) are statistically significant (t-test, $p < 0.05$).

Evaluation of the Subjective Nouns: Sentence Classification

Classifier	Acc	Prec	Rec	F1
<i>Majority Class Classifier</i>	59.0	59.0	100.0	74.2
Bag of Words	73.3	81.7	70.9	75.9
WBO	72.1	76.0	77.4	76.7
WBO+SubjNoun	74.3	78.6	77.8	78.2
WBO+SubjNoun>manual+discourse	76.1	81.3	77.4	79.3

There is a clear benefit of SubjNoun when added as a feature.

Conclusion

- ▶ Automatic extraction of subjective nouns via bootstrapping.
- ▶ Proposed algorithm requires manually designed seeds and unlabeled corpus.
- ▶ Output requires manual review but, in general, there are many subjective nouns contained in the output.
- ▶ Output of bootstrapping improves sentence-level classification.

Discussion

- ▶ Collect pros and cons for the different approaches.

References

- ▶ R. Bruce, J. Wiebe: “Recognizing subjectivity: A case study of manual tagging”, *Natural Language Engineering*, 2000.
- ▶ V. Hatzivassiloglou, K. McKeown: “Predicting the semantic orientation of adjectives”, in *ACL/EACL*, 1997.
- ▶ V. Hatzivassiloglou, J. Wiebe: “Effects of adjective orientation and gradability on sentence subjectivity”, in *COLING*, 2000.
- ▶ J. Wiebe, T. Wilson, M. Bell: “Development and use of a gold standard for subjectivity classification”, in *ACL*, 1999.