Social gender and the structure of the French derivation system

Marine Wauquier[1] Olivier Bonami[2]

[1]Université Sorbonne Nouvelle, Laboratoire Lattice

[2]Université Paris Cité, CNRS, Laboratoire de linguistique formelle

Initially presented at Gender workshop, Berlin, June 2023

The starting point

- French suffixes *-eur*$_{M}$, *-euse*$_{F}$ and *-rice*$_{F}$ derive
  - Instrument nouns:
    - *moteur* ‘motor’, *viseur* ‘(gun) sight’, …
    - *agrafeuse* ‘stapler’, *badgeuse* ‘clocking terminal’, …
    - *calculatrice* ‘pocket calculator’, *excavatrice* ‘excavator’, …
  - Agent nouns:
    - *chanteur* ‘male singer’, *monteur* ‘male film editor’, *directeur* ‘male director’, *réparateur* ‘repairman’ …
    - *chanteuse* ‘female singer’, *monteuse* ‘female film editor’, …
    - *directrice* ‘female director’, *réparatrice* ‘repair-woman’, …
- We have a clear instance of morphological rivalry:
  - For instrument nouns, the three suffixes convey the same meaning.
  - For agent nouns:
    - *-eur*$_{M}$ vs. {*-euse*$_{F}$, *-rice*$_{F}$} convey social gender information.
    - On first analysis, *-euse*$_{F}$ and *-rice*$_{F}$ convey the same meaning.

A caveat: gender *information*, not gender

- One should be careful in the description of the relationship between social and grammatical gender in languages like French.
- Vast and expanding set of morphologically related pairs of animate masculine and feminine nouns (Bonami & Boyé, 2019)

MAS	FEM	trans.
avocat	*avocate*	'lawyer'
fermier	*fermière*	'farmer'
magicien	*magicienne*	'magician'
journaliste	*journaliste*	'journalist'

MAS	FEM	trans.
chanteur	*chanteuse*	'singer'
monteur	*monteuse*	'editor'
directeur	*directrice*	'director'
réparateur	*réparatrice*	'repair p.'

- With these paired nouns:
  - Feminines pick out women.
  - Masculines can be used to pick out men or women.
  - Speakers signal gender ideology by using a masculine or feminine to refer to women (Burnett & Bonami, 2019).
  - Gender stereotypes influence the choice of gender in production (Pozniak & Burnett, 2021) and in interpretation (Gygax et al., 2012).
- This is why we talk about conveying *social gender information*.

-euse vs. *-rice*: a first pass

- Early work gave circumstantial evidence for a difference in prestige between *-euse* and *-rice* nouns (Dawes, 2003; Lenoble-Pinson, 2008).
  - *serveuse* 'waitress', *entraîneuse* 'nightclub hostess', *allumeuse* 'tease', …
  - *directrice* 'female director', *inspectrice* 'female inspector', *sénatrice* 'female senator', …
- While the examples are suggestive, these early works:
  1. do not quantify the importance of the phenomenon over the lexicon;
  2. do not rely on an independent classification of the 'prestige' associated with noun meanings.
- Wauquier et al. (2020a) addresses these concerns
  1. by examining all deverbal nouns in *-euse* (302) and *-rice* (73) in the Lexeur database (Wauquier et al., 2020b) with a frequency of at least 5 in the wikipedia corpus;
  2. by using methods from distributional semantics to assess systematically differences between the two sets of nouns.
- Their study provides a basic confirmation of partial semantic specialization.

The larger context: learnèd morphology I

- An important factor ignored by Wauquier et al. (2020a): the form of the suffix is not the only thing setting apart *-euse* and *-rice* nouns.
- Basic bifurcation in French deverbal word formation between learnèd vs. nonlearnèd processes.
  - Learnèd processes originate in vocabulary borrowed from Latin from Middle French on (Rainer & Buridant, 2015).
  - This massive influx of vocabulary led to new productive derivational processes distinctively based on a stem allomorph that is not otherwise used in the inflection system.

Verb	*-rice*	*-eur*	*-ion*	*-if*
former 'train'	*formatrice* 'female trainer'	*formateur* 'male trainer'	*formation* 'training '	*formatif* 'formative'
répéter 'repeat'	*répétitrice* 'female rehearser'	*répétiteur* 'male rehearser'	*répétition* 'rehearsal'	*répétitif* 'repetitive'
distribuer 'distribute'	*distributrice* 'female distributor'	*distributeur* 'male distributor'	*distribution* 'distribution'	*distributif* 'distributive'
ouïr 'hear'	*auditrice* 'female hearer'	*auditeur* 'male hearer'	*audition* 'hearing'	*auditif* 'auditory'

The larger context: learnèd morphology II

- By contrast, nonlearnèd processes originate in (evolutions of) inherited vocabulary.
- They normally rely on the basic stem manifest e.g. in the verb's present participle.

Verb	PRS.PTCP	*-euse*	*-eur*	*-age*	*-able*
laver 'wash'	*lavant*	*laveuse* 'female washer'	*laveur* 'male washer'	*lavage* 'washing'	*lavable* 'washable'
finir 'finish'	*finissant*	*finisseuse* 'female finisher'	*finisseur* 'male finisher'	*finissage* 'finishing'	*finissable* 'finishable'
balayer 'sweep'	*balayant*	*balayeuse* 'female sweeper'	*balayeur* 'male sweeper'	*balayage* 'sweeping'	*balayable* 'sweepable'
abattre 'take down'	*abattant*	*abatteuse* 'woman who fells'	*abatteur* 'man who fells'	*abattage* 'slaughter'	*abattable* 'slaughterable'

The larger context: learnèd morphology III

- Accordingly, there are also two processes forming masculine agent nouns, using the same affix but different stem allomorphs.

-rice	*-eur*
formatrice 'female trainer'	*formateur* 'male trainer'
répétitrice 'female rehearser'	*répétiteur* 'male rehearser'
distributrice 'female distributor'	*distributeur* 'male distributor'
auditrice 'female hearer'	*auditeur* 'male hearer'

Learnèd formations

-euse	*-eur*
laveuse 'female washer'	*laveur* 'male washer'
finisseuse 'female finisher'	*finisseur* 'male finisher'
balayeuse 'female sweeper'	*balayeur* 'male sweeper'
abatteuse 'woman who fells'	*abatteur* 'man who fells'

Nonlearnèd formations

- Note that most morphological families use only one type of formation for all their agent nouns.
- Notable but very rare exceptions with some stems in -t:
  - *enquêteur* 'male invetigator', *enquêteuse, enquêtrice* 'female investigator'
  - *sculpteur* 'male sculptor', *sculpteuse, sculptrice* 'female sculptor'

A new question

- Given that masculine agent nouns exhibit morphological contrasts parallel to those of feminine agent nouns, do they also exhibit the same semantic differences?
- Hypothesis: the relevant interpretive contrasts are associated with learnedness rather than the two suffixes *-euse* and *-rice*.

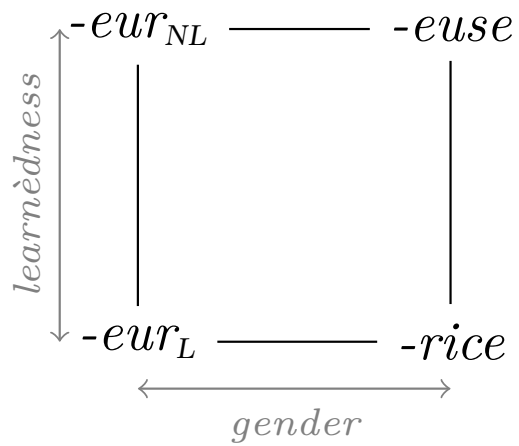

A new question

- Pushing the hypothesis further, we may find similar contrasts in other parts of the derived vocabulary, e.g. in action nouns.

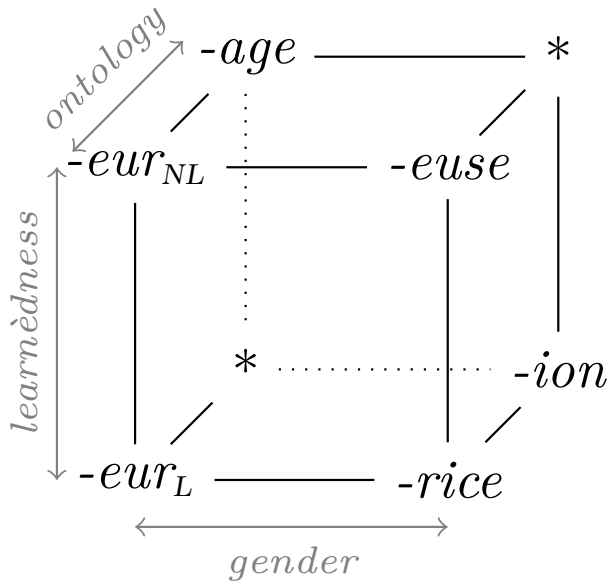

Roadmap

- In this talk we use computational methods from distributional semantics (e.g. Boleda 2020) to assess how parallel are contrasts between matching learnèd and nonlearnèd morphological categories, across three morphosemantic categories:
  - Feminine agent nouns in *-rice* vs. *euse*.
  - Learnèd vs. nonlearnèd agent nouns in *-eur*.
  - Feminine action nouns in *-ion* vs. masculine action nouns in *-age*
- Two experiments:
  1. We establish strong parallelism between the distributional properties of the three pairs of morphological categories, using a computational classification task.
  2. We show that this parallelism still implements subtly different contrasts through qualitative examination of distributional neighborhoods.

Distributional semantics in a nutshell

- Harris (1954, p. 156): “Difference of meaning correlates with difference of distribution.”
- Hence distribution can be used as a proxy for meaning.
- Modern distributional semantics relies on high-dimensional numeric vectors as representations of distribution.
- The basic intuition: vectors of cooccurrence counts.

	Contexts	
	beautiful	*thoughtful*
singer	10	1
dancer	7	3
researcher	2	10

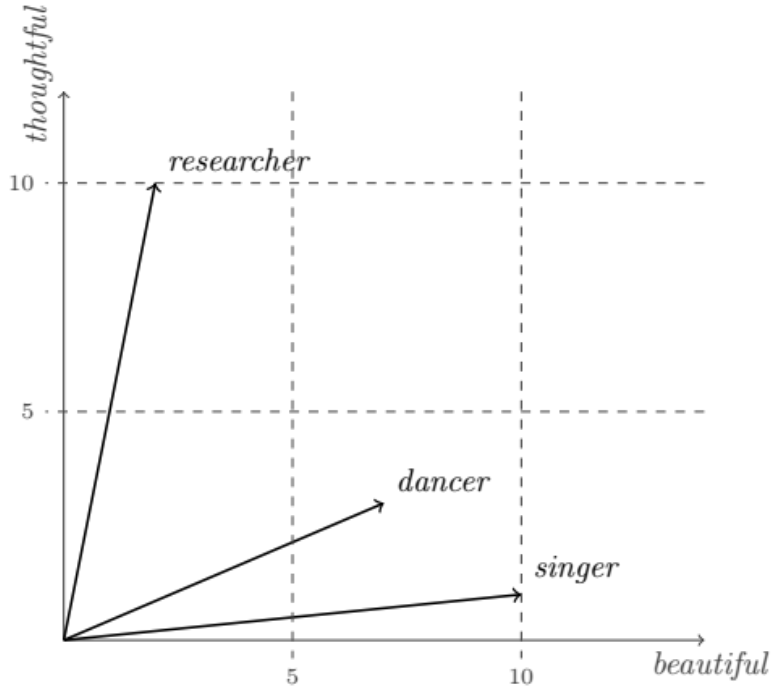

- Contemporary research uses as vectors internal states of a neural network learning to predict text.

Our vector space

- Derived from the FRCOW corpus (Schäfer, 2015; Schäfer & Bildhauer, 2012) using the Gensim (Řehůřek, 2010) implementation of word2vec (Mikolov et al., 2013).

  Hyperparameters: CBOW, 2 training epochs, 5 negative samples, window size 5, vector size 100.

- We need vectors for lexemes rather than wordforms.
- To this end we built a version of the corpus with:
  - Lemmas rather than wordforms.
    - e.g. *dînera* $\rightsquigarrow$ `dîner_ver`
  - Tagged lemmas rather than bare lemmas
    - e.g. *un dîner* $\rightsquigarrow$ `un_art dîner_nom`
  - Careful gender-neutralization
    - e.g. *du* $\rightsquigarrow$ `de_prep le_art`

  …and used that as input for word2vec.

Classification

- A (binary) classifier is an algorithm that learns to predict a (binary) distinction based on a collection of predictors.
  - Training:

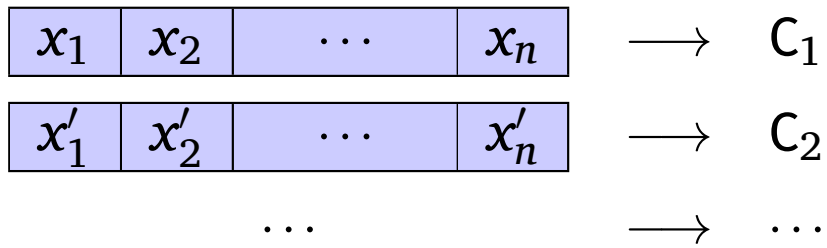

  - Testing:

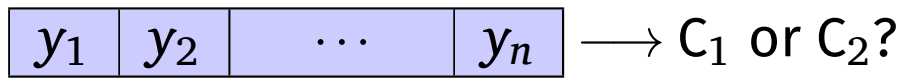

- Usually, a classifier is trained and tested on the same kind of data.
  - This is what we call intrinsic prediction.
- However, we can also test a classifier on data that is qualitatively different from the training data.

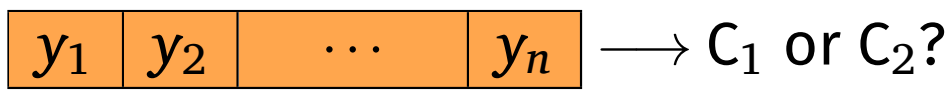

  - This is what we call extrinsic prediction.

Predictions

- In our case:
  - The predictors are the dimensions of distributional vectors.
  - The predicted distinction is learnèd vs. nonlearnèd.
- We trained classifiers to predict learnèdness in one morphosemantic category, and we test it on either the same or a different morphosemantic category.

	Test data		
Training data	FAN	MAN	ACT
Feminine Agent Nouns (FAN)	intrinsic	extrinsic	extrinsic
Masculine Agent Nouns (MAN)	extrinsic	intrinsic	extrinsic
ACtion Nouns (ACT)	extrinsic	extrinsic	intrinsic

- If learnèdness has the same distributional import in two categories, then extrinsic prediction should be just as accurate as intrinsic prediction.
- At the other extreme, if learnèdness has completely different imports in the two categories, then extrinsic prediction should be at chance level.

The data

- Sources:
  - ACTs extracted from Lexeur (Wauquier et al., 2020b)
  - FANs extracted from Lexeur with manual annotation for animacy.
  - MANs taken from the dataset of Huyghe & Wauquier (2020), with semiautomatic annotation for learnèdness.
- Dataset size:

Morphosemantic category	Learnèd	Nonlearnèd
FAN (*-rice* vs. *-euse*)	158	301
MAN (*-eur*$_{\text{L}}$ vs. *-eur*$_{\text{NL}}$)	141	462
ACT (*-ion* vs. *-age*)	750	625

  All dataset were (randomly) downsampled to 141 items, to ensure balanced comparisons.
- Classification method: gradient boosting applied to decision trees (Friedman, 2001; Mason et al., 2000).
  - Hyperparameters: 500 estimators, max depth of 2, deviance loss function.
  - 10-fold cross validation for intrinsic classification.

Classification results

Training data	Test data FAN	MAN	ACT
FAN	0.80 (0.75, 0.85)	0.77 (0.72, 0.82)	0.79 (0.74, 0.84)
MAN	0.77 (0.72, 0.82)	0.77 (0.72, 0.82)	0.82 (0.78, 0.87)
ACT	0.76 (0.71, 0.81)	0.79 (0.74, 0.84)	0.83 (0.79, 0.87)

- Intrinsic prediction is moderately accurate — but rather impressively so given the small training datased.
- No significant difference in accuracy for intrinsic prediction across morphosemantic categories.
- Crucially, no significant difference between intrinsic and extrinsic prediction.
- This supports the conclusion that, at a macroscopic level, learnèdness has the same distributional consequences across morphosemantic categories.

Discussion

- We have found strong evidence that
  - agent nouns in *-rice* and *-euse*,
  - learnèd and nonlearnèd agent nouns in *eur*,
  - action nouns in *-ion* and *-age*

  manifest the same distributional contrasts.
- This is confirmed by
  - Examination of the import of crucial vector dimensions.
  - Examination of agreement across the predictions of the different classifiers.
- However this does not show that learnèdness conveys the same meaning in all three cases.
  - Classification is a blunt instrument, and maybe there are more subtle differences to be uncovered.
  - It seems unlikely that the same differences in meaning matter for descriptions of individuals and of events.
- We hence turn to a more qualitative but more fine-grained method.

Centroids

- The centroid of a set of vectors S is the "average" vector, i.e., the vector obtained by averaging dimension by dimension the members of S.

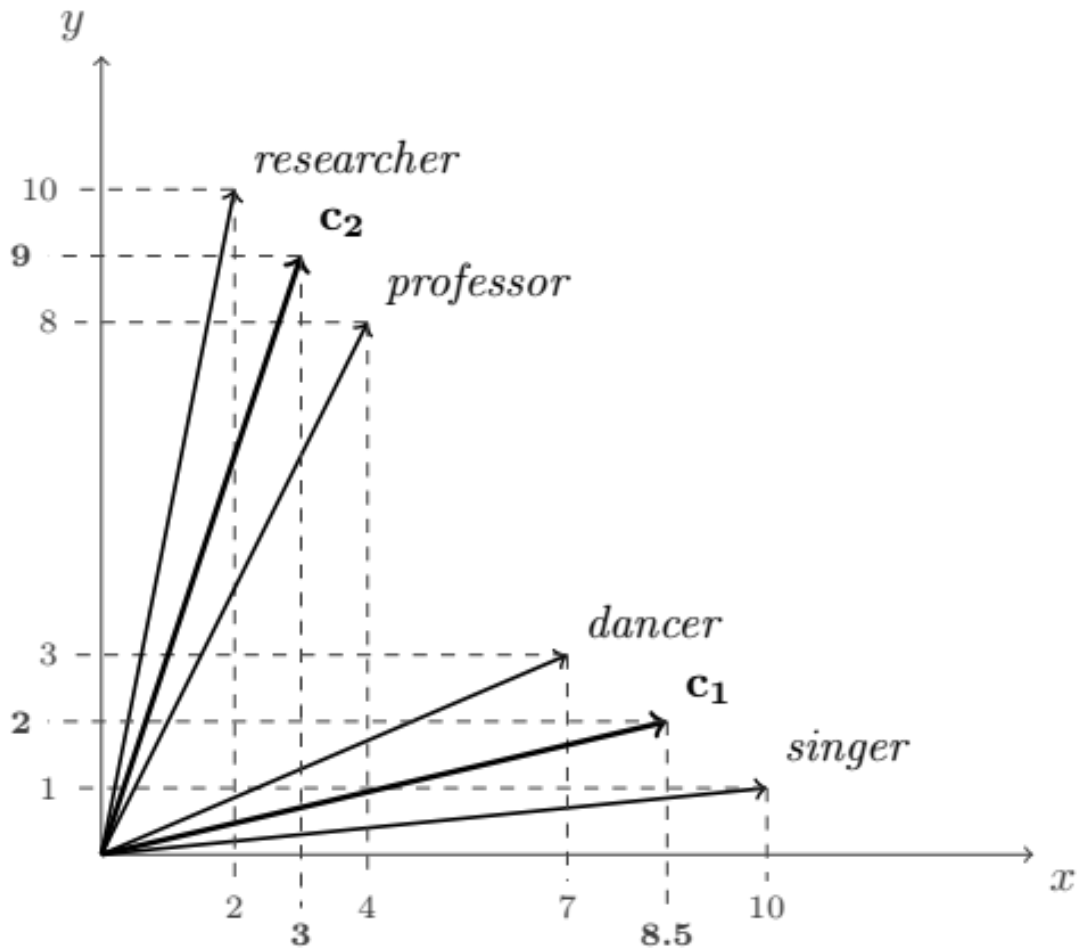

The method

- Given two morphological categories of interest, e.g. FANs in *-euse* and *-rice*:
  1. We compute the centroid of the category
     - This arguably captures what members of the category have in common
  2. We identify the 100 nearest neighbors of the centroid, using cosine similarity.
  3. We examine qualitatively the properties of these neighbors.
     - Neighbors of the centroid are semantically closest to what members of the category have in common.

-euse and *-rice*

- 5 nearest neighbors:

-euse neighbors		*-rice* neighbors	
tatoueuse	'female tattoo artist'	*comédienne*	'actress'
globetrotteuse	'female globe-trotter'	*dessinatrice*	'female artist'
laideron	'ugly woman'	*écrivaine*	'female writer'
écuyère	'horsewoman'	*programmatrice*	'female programmer'
cascadeuse	'stuntwoman'	*collaboratrice*	'female associate'

- More generally:
  - 53 of the neighbors of *-euse* have some kind of derogatory meaning.
  - These have to do with, by order of importance:
    - stigmatized sexuality, e.g. *nymphomane* 'nymphomaniac' (25)
    - negatively valued behavior, e.g. *pimbêche* 'insolent woman' (13)
    - pejorative physical characterizations, e.g. *laideron* 'ugly woman' (6)
    - etc.
  - By contrast, only 4 of the neighbors of *-rice* are derogatory. Most neighbors are either neutral or denote prestigious roles, e.g. *dirigeante* 'leader', *chirurgienne* 'surgeon', *avocate* 'lawyer'.
- This basically replicates Wauquier et al.'s (2020) results.

-eur$_{NL}$ and *-eur*$_{L}$

- 5 nearest neighbors:

-eur$_{NL}$ neighbors		*-eurŁ* neighbors	
dragueur	'womanizer'	*informateur*	'informer'
truand	'gangster'	*exécutant*	'subordinate'
armurier	'gunsmith'	*commanditaire*	'silent partner'
artificier	'pyrotechnist'	*savant*	'scholar'
tâcheron	'drudge'	*pédagogue*	'educator'

- More generally:
  - 30 of the neighbors of *-eur*$_{NL}$ have some kind of derogatory meaning.
  - These have to do with:
    - criminality, e.g. *truand* 'gangster' (11)
    - all kinds of negatively valued behaviours, e.g. *poivrot* 'drunkard' (16)
    - Much more rarely, sexuality, e.g. : *dragueur* 'womanizer' (2)
  - By contrast, only 8 of the neighbors of *-eur*$_{L}$ are derogatory. Most neighbors are either neutral or denote prestigious roles, e.g. *érudit* 'scholar', *académicien* 'academician', *orateur* 'orator'.
- Interestingly, we see a similar contrast, but not building on the same stigmatized characteristics.

-age and *-ion*

- 5 nearest neighbors:

-age neighbors		*-ion* neighbors	
compactage	'compacting'	*assimilation*	'assimilation'
ponçage	'sanding'	*dissociation*	'dissociation'
sablage	'sandblasting'	*dénaturation*	'denaturation'
meulage	'grinding'	*disjonction*	'disjunction'
piquage	'pricking'	*transformation*	'transformation'

- More generally:
  - Neighbors of *-age* largely name industrial or manufacturing processes (*meulage* 'grinding', *ragréage* 'screeding') or actions of lower complexity (*nettoyage* 'cleaning', *clouage* 'nailing').
  - Neighbors of *-ion* name actions pertaining to scientific domains (*cristallisation* 'crystallization', *immunosuppression* 'immunosuppression') or more abstract concepts (*hiérarchisation* 'hierarchization', *généralisation* 'generalization')
- This replicates the results of Wauquier et al. (2020a).

Discussion

- With agent nouns, learnèd vs. nonlearnèd morphology codes axiological valence: nonlearnèd nouns tend to be negatively valued.
- However, the relevant axiological properties are not the same in the masculine and in the feminine.
- This is likely a reflex of gender stereotypes: negative properties of men and women are different.
- Interesting connection with sociological work on feminities and masculinities:
  - Discourses about women tend to stigmatize (i) immoral and (ii) infantile or incompetent behavior (Bosmajian, 1977).
  - Criminality plays a special role in the characterization of masculinities. (Messerschmidt, 1993; Connell, 1995).
- Learnèd vs. nonlearnèd morphology codes strikingly different semantic contrasts for agent nouns vs. action nouns: prestige of activities vs. axiological valence.

Conclusions

- Using distributional semantics allowed us to document subtle semantic consequences of learnèdness distinctions in the French derived lexicon.
  - We suspect that similar effects are likely to be found in all situations where sources of morphology in a language differ in social status.
- We confirmed the existence of a prestige distinction among processes forming feminine agent nouns, but documented a parallel distinction in masculine agent nouns.
- We showed that the basic morphosemantic distinction interacts in a subtle way with gender stereotypes.
- More generally, our method provides a new kind of evidence to explore gender ideology and its interaction with language.

Thanks

ANR Project *Demonext*