Voronoi Languages

Gerhard Jäger
gerhard.jaeger@uni-tuebingen.de
joint work with Lars Metzger and Frank Riedel

May 28, 2010

Workshop on *Game Theory and Communication*, Stanford

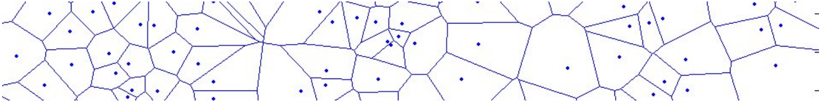

Eberhard Karls Universität Tübingen

sfs seminar für sprachwissenschaft

Overview

- Signaling games with a Euclidean meaning space: the model
- structure of Nash equilibria
- evolution: finite strategy space
- evolution: infinite strategy space
- applications and modifications

Signaling game

- two players:
  - **S**ender
  - **R**eceiver
- set of **M**eanings
- finite set of **F**orms
- sequential game:
  1. nature picks out $m \in M$ according to some probability distribution p and reveals m to S
  2. S maps m to a form f and reveals f to R
  3. R maps f to a meaning m'

Signaling game

- standard utility function (extensive form):

$$u_{s/r}(m, f, m') = \begin{cases} 1 & \text{if } m = m' \\ 0 & \text{else} \end{cases}$$

or perhaps

$$u_{s/r}(m, f, m') = -\text{cost}(f) + \begin{cases} 1 & \text{if } m = m' \\ 0 & \text{else} \end{cases}$$

Euclidean meaning space

- Modification of standard model:
  - graded notion of **similarity** between meanings
  - players try to maximize similarity between m and m'
  - implementation using *conceptual spaces*:
    - meanings are points in n-dimensional Euclidean space
    - similarity is inversely related to **distance**
  - large set of meanings, small set of forms
- Linguistic motivation:
  - lexical semantics, esp. of simple adjectives
  - finite categorization of continuous high-dimensional space
  - possible connections to cognitive psychology and quantitative distributional semantics

Utility function

General format

$$u_{s/r}(m, f, m') \quad = \quad \text{sim}(m, m')$$

- $\text{sim}(x, y)$ is strictly monotonically decreasing in Euclidean distance $\|x - y\|$

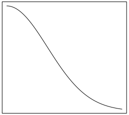

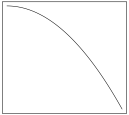

in this talk, I assume either

- a **Gaussian** similarity function

  $$\text{sim}(x, y) \doteq \exp(-\frac{\|x - y\|^2}{2\sigma})$$

  (psychologically plausible), or

- a **quadratic** dependency

  $$\text{sim}(x, y) \doteq -\|x - y\|^2$$

  (better mathematical tractability)

Normal form

- prior probability density function f over meanings ("nature") is exogenously given
- set of meanings is a **finite** or a **convex and compact** subset of $\mathbb{R}^n$
- normalized utility functions (S and R are sender/receiver strategies resp.)

Finite meaning space

$$u_{s/r}(S, R) = \sum_m f(m)\mathrm{sim}(m, R(S(m)))$$

Continuous meaning space

$$u_{s/f} = \int_{\mathbb{R}^n} f(x)\mathrm{sim}(x, R(S(x)))dx$$

Evolution of strategies

- main interest of this talk: which strategy pairs are **dynamically stable** under evolution?
- evolutionary dynamics:
  - replicator dynamics
  - utility = replicative success
  - idealizations:
    - infinite population
    - everybody interacts with everybody else with equal probability
- dynamic stability concepts

  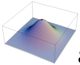

  asymptotically stable point: dynamically attracts all points that are *sufficiently close* (according to some suitable notion of distance between population states)

  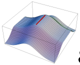

  asymptotically stable set: continous (compact) set of points that jointly attract all points that are outside the set but sufficiently close

Simulations

- two-dimensional circular meaning space
- finitely many pixels (meanings)
- uniform distribution over meanings
- initial stratgies are randomized
- update rule according to (discrete time version of) replicator dynamics

Simulations

- two-dimensional circular meaning space
- finitely many pixels (meanings)
- uniform distribution over meanings
- initial stratgies are randomized
- update rule according to (discrete time version of) replicator dynamics

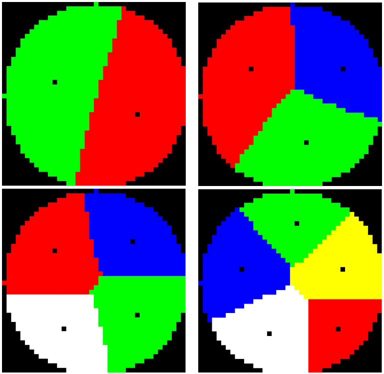

Voronoi tesselations

- suppose R (a pure strategy) is known to the sender: which sender strategy would be the best response to it?
  - every form f has a “prototypical” interpretation: $R(f)$
  - for every meaning m: S’s best choice is to choose the f that minimizes the distance between m and $R(f)$
  - optimal S thus induces a **partition** of the meaning space
  - Voronoi tesselation, induced by the range of R
  - tiles in a Voronoi tesselation are always **convex**

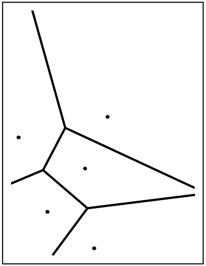

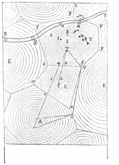

Nash equilibria

- suppose S (also pure) is known to the receiver: which receiver strategy is a best response?
  - receiver has map each signal f to a point that maximizes average similarity to the points in $S^{-f}(f)$
  - intuitively, this is the center of f's Voronoi cell
  - formally: if R is a best response to S, then

$$R(f) = \arg_x \min \int_{S^{-1}(f)} f(y)\mathrm{sim}(x, y)dy$$

- for continuous meaning space always uniquely defined
- for a quadratic similarity function, this is the center of gravity of the Voronoi cell:

$$R(f) = \int_{S^{-1}(f)} f(y)ydy$$

Evolutionary stability in finite strategy space: static notion

Theorem (Selten 1980)

In asymmetric games, the evolutionarily stable states are exactly the strict Nash equilibria.

- In asymmetric games and in partnership games, the asymptotically stable states are exactly the ESSs (Cressman 2003; Hofbauer and Sigmund 1998)
- asymptically stable state entails Voronoi tesselation
- This does **not** entail (yet) that evolution always leads to Voronoi strategies

Evolutionarily stable sets

- some games do not have an ESS

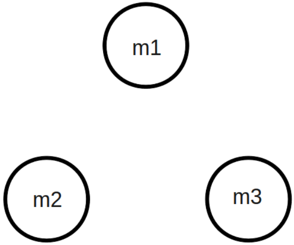

- evolution nevertheless leads to Voronoi languages

Evolutionary stability in finite strategy space: static notion

Definition

A set E of symmetric Nash equilibria is an *evolutionarily stable set* (ESSet) if, for all $x^* \in E, u(x^*, y) > u(y, y)$ whenever $u(y, x^*) = u(x^*, x^*)$ and $y \notin E$. (Cressman 2003)

Observation

If R is a pure receiver strategy, the inverse image of any $S \in BR(R)$ is consistent with the Voronoi tessellation of the meaning space that is induced by the image of R.

Evolutionary stability in finite strategy space: static notion

Theorem

If a symmetric strategy is an element of some ESSet, the inverse image of its sender strategy is consistent with the Voronoi tessellation that is induced by the image of its receiver strategy.

sketch of proof:

- game in question is symmetrized asymmetric game
- ESSets of symmetrized games coincide with SESets of asymmetric game (Cressman, 2003)
- SESets are sets of NE
- SESets are finite unions of Cartesian producs of faces of the state space
- hence every component of an element of an SESet is a best reply to some pure strategy

Static and dynamic stability in finite strategy space

Asymptotic stability

- in symmetrized games with a finite strategy space, a set E is an *asymptotically stable set of rest points* if and only if it is an ESSet
- in partnership games, at least one ESSet exists
- intuitive interpretation: under replicator dynamics $+$ small effects of drift, system will eventually converge into some ESSet

Dynamic stability in games with continuous strategy spaces

- in finite games, every strict Nash equilibrium is asymptotically stable
- for games with a continuum of strategies, things are more complex ... (cf. for instance Oechssler and Riedel 2001)
  - definition of stability refers to *topology* of the state space, i.e. to a notion of *closeness* between population states
  - population state: probability measure over strategies
  - finite strategy space: closeness of states means closeness of probabilities for each strategy
  - continuous strategy space: small deviation means
    - few agents change their strategy drastically, or
    - many agents change their strategy slightly
  - every asymptotically stable point (set) is an ESS (ESSet), but not vice versa

Dynamic stability in games with continuous strategy spaces

Example

$$u(x, y) = -x^2 + 4xy$$

all real numbers are possible strategies

- $(0, 0)$ is a strict Nash equilibrium
- homogeneous 0-population cannot be invaded by a single mutant with a different strategy
- if entire population mutates to some $\epsilon \neq 0$, it will **not** return to the equilibrium, no matter how small $|\epsilon|$ is

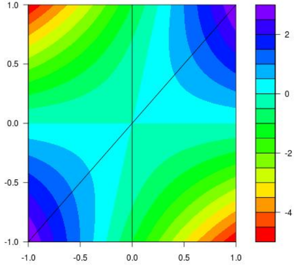

Signaling games with continuous meaning space

- each such game has an asymptotically stable rest point
- sketch of proof:
  - in partnership games, utility is a Lyapunov function
  - utility is continuous is state space
  - state space is compact
  - hence utility has a maximum, which must then be asymptotically stable
- every trajectory converges to some as. st. state
- all asymptotically stable states are strict Nash equilibria
- as in previous example, not every strict NE is as. st.
- several static stability notions have been suggested in the literature, but none coincides with dynamic stability for the class of games considered here

Signaling games with continuous meaning space

Example

- meaning space: unit square $[0,1] \times [0,1]$
- uniform probability distribution
- quadratic similarity function
- two signals

Signaling games with continuous meaning space

Example

- meaning space: unit square $[0, 1] \times [0, 1]$
- uniform probability distribution
- quadratic similarity function
- two signals

- two strict Nash equilibria (up to symmetries)

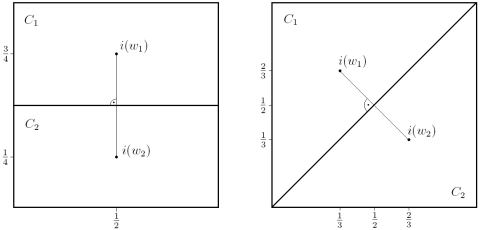

- only the left one is dynamically stable

Stability vs. efficiency

Example

- meaning space: rectangle $[0, a] \times [0, b]$ with $3b^2 > 2a^2$
- uniform probability distribution
- quadratic similarity function
- two signals

Stability vs. efficiency

Example

- meaning space: rectangle $[0, a] \times [0, b]$ with $3b^2 > 2a^2$
- uniform probability distribution
- quadratic similarity function
- two signals

- two dynamically stable states

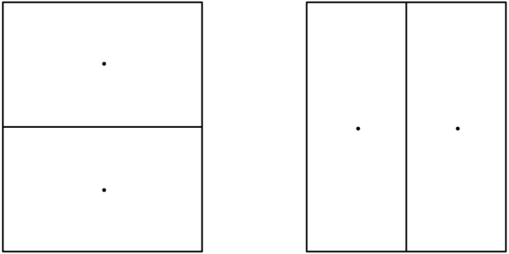

- the left one has a higher utility than the right one
- this means that the left equilibrium is sub-optimal but nevertheless stable

Unit square, three words

Unit square, three words

- four strict equilibria (up to symmetries)

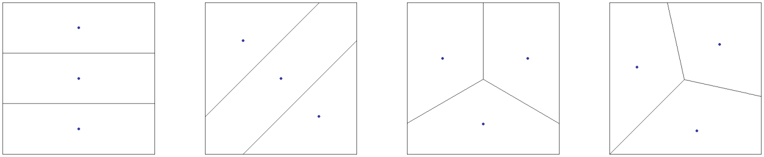

Unit square, three words

- four strict equilibria (up to symmetries)

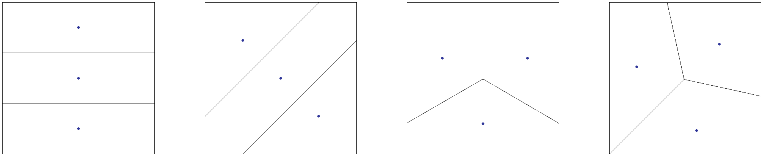

- only the first one is dynamically stable

Unit square, many words

- for small number of words, square shaped cells are stable
- for larger numbers, evolution favors hexagonal cells

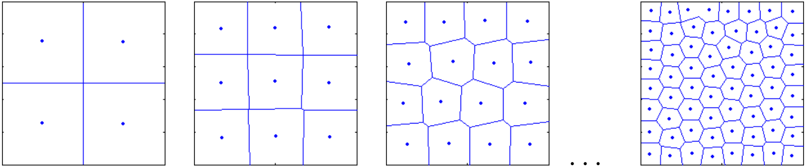

. . .

Skewed probability distributions

- uniform probability distribution over meanings favors tesselation into regular polygons
- skewed distributions lead to irregular shapes
- tendency: high probability regions are covered by small tiles
- no analytical results about this so far though

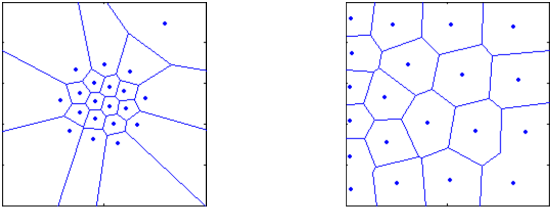

Potential application: color categorization

The color solid

- psychological color space
  - three-dimensional
  - Euclidean topology (where distances reflect subjective similarities)
  - irregularly shaped spindle-like object


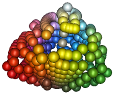

The Munsell chart

- 2d-rendering of the surface of the color solid
  - 8 levels of lightness
  - 40 hues
- plus: black–white axis with 8 shaded of grey in between
- neighboring chips differ in the minimally perceivable way

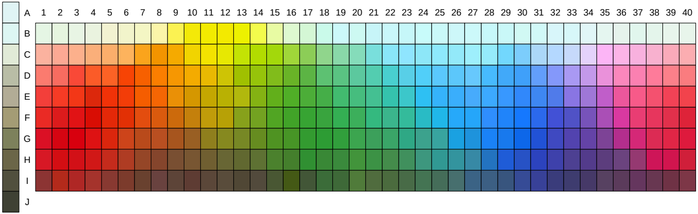

The World Color Survey

- building on work by Berlin and Kay, in 1976 Kay and co-workers launched the *world color survey*
- investigation of 110 non-written languages from around the world
- around 25 informants per language
- two tasks:
  - the 330 Munsell chips were presented to each test person one after the other in random order; they had to assign each chip to some basic color term from their native language
  - for each native basic color term, each informant identified the prototypical instance(s)

Convex color categories

- categorization task leads to partition of Munsell space for each participant
- raw data are noisy; statistical dimensionality reduction yields smooth partitions (cf. Jäger 2009; Jäger 2010)

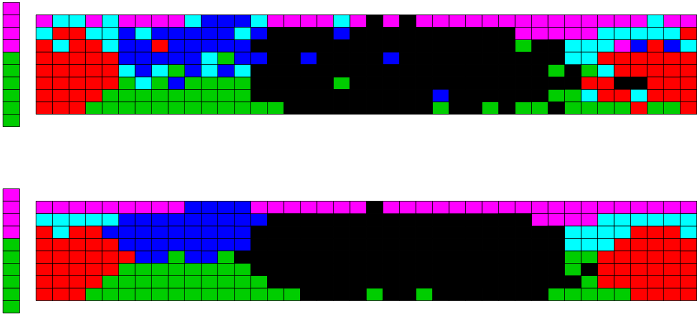

raw and processed data from a randomly picked WCS participant

Convex approximation

- on average, 93.7% of all Munsell chips are correctly classified by best convex approximation
- only small number of possible tesselations (up to some minor variation)[a]
- question for future research: are these partitions Voronoi?
- if so: can we somehow estimate the prior probabilities for colors to come up with actual empirical (here: typological) predictions?

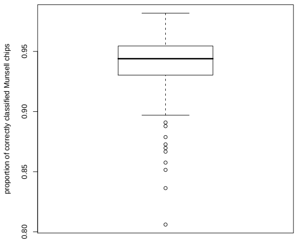

[a] Things are not quite as clear-cut as Berlin and Kay would have it though.

Finitely many meanings, continuous signal space

Related modification of standard model

- finitely many meanings
- continuum of forms (points in a Euclidean space)
- noisy transmission
- noise is normally distributed

Signaling with noisy transmission

Strict Nash equilibria

- sender strategy: mapping from vowel categories to points in the meaning space
- receiver strategy: categorization of signals

Voronoi tesselations

- suppose receiver strategy R is given and known to the sender: which sender strategy would be the best response to it?
  - every signal f has a "prototypical" interpretation: $R(f)$
  - for every meaning m: S's best choice is to choose the f that minimizes the distance between m and $R(f)$
  - optimal S thus induces a Voronoi tesselation of the signal space

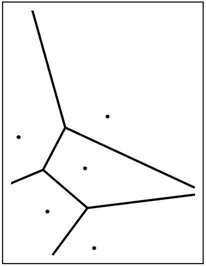

Extreme prototypes

Best response of the sender

- suppose strategy of receiver — essentially a partition of the signal space — is known
- best response of the sender is to **maximize** distance to boundaries of this partition
- if a partition cell is at the boundary of the signal space, the prototype is not central but extreme

Application: vowel space

- meanings: vowel phonemes
- signals: points in acoustic F1/F2 space

Simulations

- colored dots display receiver strategies

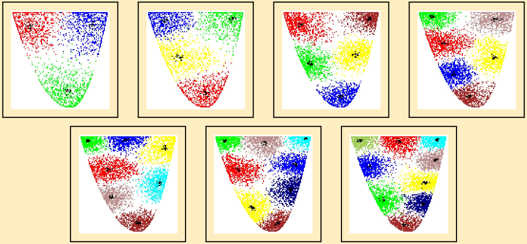

- suggestive similarity to typologically attested patterns

Directions for future work

- more specific generalizations on relation probability distribution/equilibrium structure
- impact of costs
- same question for game with noisy signals
- endogenization of prior probabilities
- other metrical spaces
- non-partnership games

Berlin, B. and P. Kay (1969). *Basic color terms: their universality and evolution*. University of California Press, Chicago.

Cressman, R. (2003). *Evolutionary Dynamics and Extensive Form Games*. MIT Press, Cambridge, Mass.

Hofbauer, J. and K. Sigmund (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK.

Jäger, G. (2009). Natural color categories are convex sets. In M. Aloni, H. Bastiaanse, T. de Jager, P. van Ormondt, and K. Schulz, eds., *Seventeenth Amsterdam Colloquium. Pre-proceedings*, pp. 11–20. University of Amsterdam.

Jäger, G. (2010). Using statistics for cross-linguistic semantics: a quantitative investigation of the typology of color naming systems. ms., submitted to *Journal of Semantics*.

Oechssler, J. and F. Riedel (2001). Evolutionary dynamics on infinite strategy spaces. *Economic Theory*, **17**(1):141–162.

Selten, R. (1980). A note on evolutionarily stable strategies in asymmetric animal conflicts. *Journal of Theoretical Biology*, **84**:93–101.