

Introducing Hatebase: the world's largest online database of hate speech The Sentinel Project for Genocide Prevention

25 March 2013

Predicting genocide is, by definition, an almost impossible task due to the scarcity of early, actionable data. There's no chi-squared test or Monte Carlo method for reliably distributing societies along a spectrum from homogeneous to homicidal, both because the extermination of entire populations has become a relatively rare occurrence (thanks to the ever-increasing internationalization of human rights, law, media, and trade) and because those societies which do succeed at systematized annihilation are often equally resourceful at hiding evidence of their crimes. (...)

Our second strategy has been to improve the tools with which we parse and prioritize data, whether from the field, from mainstream media or from social networks. To this end, the Sentinel Project recently partnered with my own organization, Mobocracy, on the development of Hatebase, an authoritative, multilingual, usage-based repository of structured hate speech which data-driven NGOs can use to better contextualize conversations from known conflict zones.

Hatebase is available to casual users through a Wikipedia-like web interface, and to developers through an authenticating API. Although the core of Hatebase is its community-edited vocabulary of multilingual hate speech, a critical concept in Hatebase is regionality: users can associate hate speech with geography, thus building a parallel dataset of "sightings" which can be monitored for frequency, localization, migration, and transformation.

For instance, an organization monitoring several simultaneous theaters of operation might integrate location-based Hatebase data into its monitoring software to assign additional real-time "weight" to specific conflict zones, providing guidance on how to best redeploy limited resources. For genocide monitoring organizations in particular, regional hate speech is a widely recognized indicator of elevated risk.

There are some weaknesses implicit in a solely vocabulary-based approach to linguistic analysis. Innocuous language, when localized, can adopt a sinister secondary meaning (e.g. "cockroaches," meaning Tutsis in Rwanda), and threats can be communicated without the need for easily identified keywords ("their days are numbered"). Despite these limitations, Hatebase can provide a layer of relevance which complements other context-based information sources, not unlike traffic congestion layered onto a city map.

In the months ahead, we'll be adding additional data attributes, visualizations, and end-user functionality to Hatebase, with a particular focus on strengthening the API in accordance with our commitment to partnership-based innovation. Our hope is that other individuals, groups and

organizations will embrace this collaborative model by leveraging Hatebase data in their own applications.

See the full article
[here.](#)