The Hegselmann-Krause model of opinion dynamics

Peter Hegarty
(with the help of: Edvin Wedin, Anders Martinsson, Mattias Danielsson, Jimmie Ekström, Jesper Johansson and Gustav Karlsson)

Department of Mathematics, Chalmers/Gothenburg University

Talk at Univeristy of Bristol, 6 Dec. 2016

- [WH1, 2015] E. Wedin and P. Hegarty, *The Hegselmann-Krause dynamics for continuous agents and a regular opinion function do not always lead to consensus*, IEEE Trans. Automat. Control **60** (2915), no. 9, 2416–2421.
- [WH2, 2015] E. Wedin and P. Hegarty, *A quadratic lower bound for the convergence rate in the one-dimensional Hegselmann-Krause bounded confidence dynamics*, Discrete Comput. Geom. **53** (2015), no. 2, 478–486.
- [HW, 2016] P. Hegarty and E. Wedin, *The Hegselmann-Krause dynamics for equally spaced agents*, J. Difference Equ. Appl. (2016). Available online.
- [HMW, 2016] P. Hegarty, A. Martinsson and E. Wedin, *The Hegselmann-Krause dynamics on the circle converge*. J. Difference Equ. Appl. (2016). Available online.

- [M, 2015] A. Martinsson, *An improved energy argument for the Hegselmann-Krause model*, J. Difference Equ. Appl. **22**, (2016), no. 4, 630–635.
- [DEJK, 2015] M. Danielsson, J. Ekström, G. Karlsson and J. Johansson, *The Hegselmann-Krause model of opinion dynamics in one and two dimensions: Phase transitions, periodicity and other phenomena*. Bachelor's Thesis, Chalmers University of Technology (2015). Available on request.

- Graph G, whose nodes are the opinionated agents.

- Graph G, whose nodes are the opinionated agents.
- Opinions are real numbers in $[0, 1]$.

- Graph G, whose nodes are the opinionated agents.
- Opinions are real numbers in $[0, 1]$.
- At discrete time steps, a random pair of neighbors in G "meet and discuss".

- Graph G, whose nodes are the opinionated agents.
- Opinions are real numbers in $[0, 1]$.
- At discrete time steps, a random pair of neighbors in G "meet and discuss".
- Parameters $\theta \in [0, 1]$ and $\mu \in (0, 1/2]$ such that, if agents with opinions (a, b) meet, then afterwards their opinions will have changed to

- Graph G, whose nodes are the opinionated agents.
- Opinions are real numbers in $[0, 1]$.
- At discrete time steps, a random pair of neighbors in G "meet and discuss".
- Parameters $\theta \in [0, 1]$ and $\mu \in (0, 1/2]$ such that, if agents with opinions (a, b) meet, then afterwards their opinions will have changed to

$$\begin{cases} (a + \mu(b - a),\ b - \mu(b - a)), & \text{if } |a - b| \le \theta, \\ (a,\ b), & \text{if } |a - b| > \theta. \end{cases}$$

- Graph G, whose nodes are the opinionated agents.
- Opinions are real numbers in $[0, 1]$.
- At discrete time steps, a random pair of neighbors in G "meet and discuss".
- Parameters $\theta \in [0, 1]$ and $\mu \in (0, 1/2]$ such that, if agents with opinions (a, b) meet, then afterwards their opinions will have changed to

$$\begin{cases} (a + \mu(b-a),\ b - \mu(b-a)), & \text{if } |a-b| \leq \theta, \\ (a,\ b), & \text{if } |a-b| > \theta. \end{cases}$$

A beautiful result is that, if $G = \mathbb{Z}$ and initial opinions are i.i.d. in $[0, 1]$, then

- Graph G, whose nodes are the opinionated agents.
- Opinions are real numbers in $[0, 1]$.
- At discrete time steps, a random pair of neighbors in G "meet and discuss".
- Parameters $\theta \in [0, 1]$ and $\mu \in (0, 1/2]$ such that, if agents with opinions (a, b) meet, then afterwards their opinions will have changed to

$$\begin{cases} (a + \mu(b - a),\ b - \mu(b - a)), & \text{if } |a - b| \leq \theta, \\ (a,\ b), & \text{if } |a - b| > \theta. \end{cases}$$

A beautiful result is that, if $G = \mathbb{Z}$ and initial opinions are i.i.d. in $[0, 1]$, then
(i) If $\theta > 1/2$ then, for any μ, almost surely all opinions converge to $1/2$.

- Graph G, whose nodes are the opinionated agents.
- Opinions are real numbers in $[0, 1]$.
- At discrete time steps, a random pair of neighbors in G "meet and discuss".
- Parameters $\theta \in [0, 1]$ and $\mu \in (0, 1/2]$ such that, if agents with opinions (a, b) meet, then afterwards their opinions will have changed to

$$\begin{cases} (a + \mu(b - a),\ b - \mu(b - a)), & \text{if } |a - b| \leq \theta, \\ (a,\ b), & \text{if } |a - b| > \theta. \end{cases}$$

A beautiful result is that, if $G = \mathbb{Z}$ and initial opinions are i.i.d. in $[0, 1]$, then
(i) If $\theta > 1/2$ then, for any μ, almost surely all opinions converge to $1/2$.
(ii) If $\theta < 1/2$ then, for any μ, almost surely disagreement persists.

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.
- Time is discrete: $t = 0, 1, \ldots$

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.
- Time is discrete: $t = 0, 1, \ldots$
- A real number $x_i(t)$ represents the opinion of agent i at time t.

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.
- Time is discrete: $t = 0, 1, \ldots$
- A real number $x_i(t)$ represents the opinion of agent i at time t.
- There is a *confidence bound* $r > 0$, which is the same for all agents.

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.
- Time is discrete: $t = 0, 1, \ldots$
- A real number $x_i(t)$ represents the opinion of agent i at time t.
- There is a *confidence bound* $r > 0$, which is the same for all agents.
- Opinions are updated synchronously according to

$$x_i(t+1) = \frac{1}{|\mathcal{N}_i(t)|} \sum_{j \in \mathcal{N}_i(t)} x_j(t),$$

where

$$\mathcal{N}_i(t) = \{j : ||x_j(t) - x_i(t)|| \leq r\}.$$

- Finite number, n say, of agents, indexed by the integers $1, 2, \ldots, n$.
- Time is discrete: $t = 0, 1, \ldots$
- A real number $x_i(t)$ represents the opinion of agent i at time t.
- There is a *confidence bound* $r > 0$, which is the same for all agents.
- Opinions are updated synchronously according to

$$x_i(t+1) = \frac{1}{|\mathcal{N}_i(t)|} \sum_{j \in \mathcal{N}_i(t)} x_j(t),$$

  where

$$\mathcal{N}_i(t) = \{j : ||x_j(t) - x_i(t)|| \leq r\}.$$

- The dynamics are unaffected by rescaling (update rule is linear), so WLOG $r = 1$.

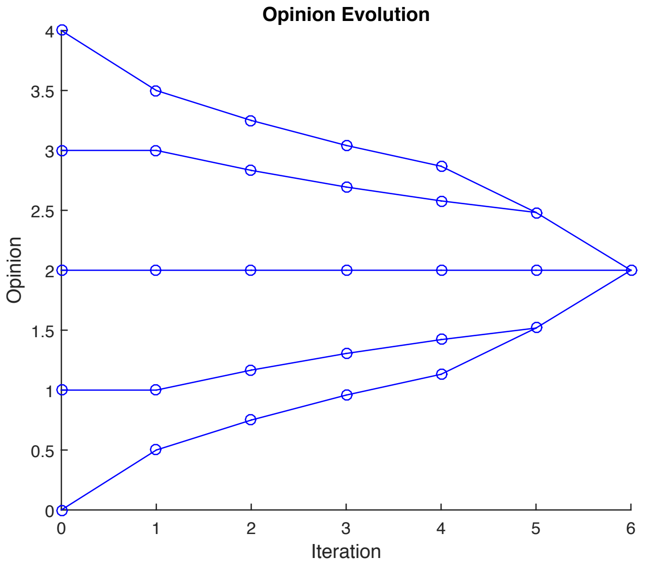
Opinion Evolution
Opinion
Iteration
4
3.5
3
2.5
2
1.5
1
0.5
0
0 1 2 3 4 5 6

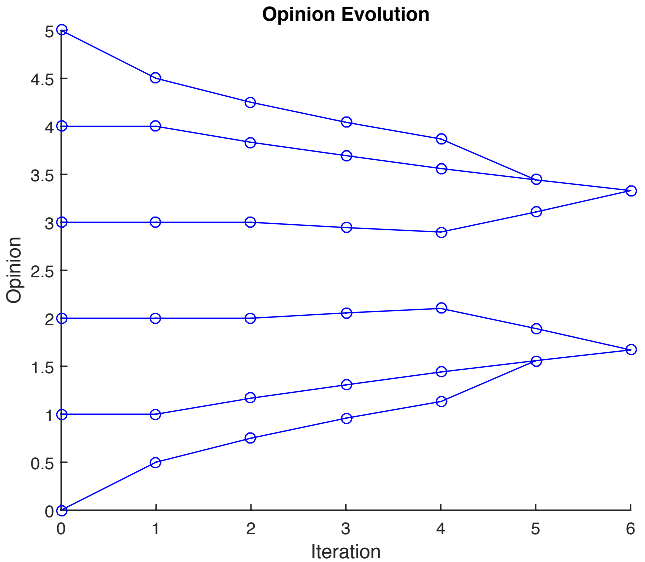
Opinion Evolution
Opinion
Iteration
0
0.5
1
1.5
2
2.5
3
3.5
4
4.5
5
0
1
2
3
4
5
6

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

Interpretation: There are k issues, and two agents must be close on all issues for compromise to occur.

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

Interpretation: There are k issues, and two agents must be close on all issues for compromise to occur. Note that a priori no reason to favour L^2-norm over any other in this interpretation.

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

Interpretation: There are k issues, and two agents must be close on all issues for compromise to occur. Note that a priori no reason to favour L^2-norm over any other in this interpretation. However, it gives the most natural geometrical interpretation, and is the one used in robotics applications (multi-agent rendezvous).

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

Interpretation: There are k issues, and two agents must be close on all issues for compromise to occur. Note that a priori no reason to favour L^2-norm over any other in this interpretation. However, it gives the most natural geometrical interpretation, and is the one used in robotics applications (multi-agent rendezvous).

Example 2. The circle $V = \mathbb{T}^1$, of diameter greater than 2.

The model makes sense if opinions are drawn from any set V with enough structure to make sense of the command to:

"Move to the average of a finite collection of points within distance one of your present location."

Example 1. Higher dimensional Euclidean space $V = \mathbb{R}^k$.

Interpretation: There are k issues, and two agents must be close on all issues for compromise to occur. Note that a priori no reason to favour L^2-norm over any other in this interpretation. However, it gives the most natural geometrical interpretation, and is the one used in robotics applications (multi-agent rendezvous).

Example 2. The circle $V = \mathbb{T}^1$, of diameter greater than 2.

Interpretation: Imagine, for example, that the issue under discussion is the time of day or year for holding some event.

Convergence in $\mathbb{R}$:

Convergence in $\mathbb{R}$:

- Very easy to show that opinions converge to limiting values (general nonsense, Banach Fixed Point Theorem blah blah ...)

Convergence in $\mathbb{R}$:

- Very easy to show that opinions converge to limiting values (general nonsense, Banach Fixed Point Theorem blah blah ...)
- In fact quite easy to show that opinions **freeze**, i.e.: there is always some $T > 0$ such that $x_i(t) = x_i(T)$ for all i and all $t \geq T$.

Convergence in $\mathbb{R}$:

- Very easy to show that opinions converge to limiting values (general nonsense, Banach Fixed Point Theorem blah blah ...)
- In fact quite easy to show that opinions **freeze**, i.e.: there is always some $T > 0$ such that $x_i(t) = x_i(T)$ for all i and all $t \geq T$.
- Still quite easy to show that the freezing time is bounded by a universal polynomial function of the number of agents:

Convergence in $\mathbb{R}$:

- Very easy to show that opinions converge to limiting values (general nonsense, Banach Fixed Point Theorem blah blah ...)
- In fact quite easy to show that opinions **freeze**, i.e.: there is always some $T > 0$ such that $x_i(t) = x_i(T)$ for all i and all $t \geq T$.
- Still quite easy to show that the freezing time is bounded by a universal polynomial function of the number of agents:
  $\Rightarrow$ Can get a bound of around $O(n^5)$ from general theory of Markov chains on graphs.

Convergence in $\mathbb{R}$:

- Very easy to show that opinions converge to limiting values (general nonsense, Banach Fixed Point Theorem blah blah ...)
- In fact quite easy to show that opinions **freeze**, i.e.: there is always some $T > 0$ such that $x_i(t) = x_i(T)$ for all i and all $t \geq T$.
- Still quite easy to show that the freezing time is bounded by a universal polynomial function of the number of agents:
  $\Rightarrow$ Can get a bound of around $O(n^5)$ from general theory of Markov chains on graphs.
  $\Rightarrow$ Best to date is $O(n^3)$. Elementary argument which considers the behaviour of the extremal agents (Bhattachrya et al, 2013).

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW].

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW].

- Easy to see that n agents placed distance one apart will take time $\Omega(n)$ to freeze.

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW].

- Easy to see that n agents placed distance one apart will take time $\Omega(n)$ to freeze.
  We proved [HW] that this configuration evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW].

- Easy to see that n agents placed distance one apart will take time $\Omega(n)$ to freeze.
  We proved [HW] that this configuration evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.
  In particular, the freezing time is $5n/6 + O(1)$.

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW].

- Easy to see that n agents placed distance one apart will take time $\Omega(n)$ to freeze.
  We proved [HW] that this configuration evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.
  In particular, the freezing time is $5n/6 + O(1)$.
- We were surprised to discover a configuration which takes time $\Omega(n^2)$ to freeze: **Dumbbell graph**

Lower bounds on a universal freezing time first studied in any seriousness by Wedin and myself [WH2, HW].

- Easy to see that n agents placed distance one apart will take time $\Omega(n)$ to freeze.
  We proved [HW] that this configuration evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.
  In particular, the freezing time is $5n/6 + O(1)$.
- We were surprised to discover a configuration which takes time $\Omega(n^2)$ to freeze: **Dumbbell graph**
- We believe that the freezing time is always $O(n^2)$, but this remains open.

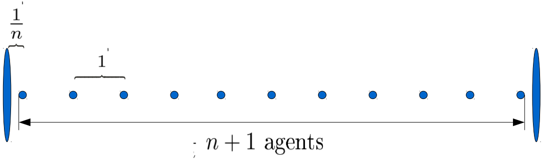

Figure: Schematic representation of the configuration $\mathcal{D}_n$. Each dumbbell has weight n.

Convergence in $\mathbb{R}^k$, $k > 1$:

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.
- Method of Bhattachrya et al does not generalize.

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.
- Method of Bhattachrya et al does not generalize.
  Moral: The jump from one to two dimensions is key.

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.
- Method of Bhattachrya et al does not generalize.
  Moral: The jump from one to two dimensions is key.
- Instead, state-of-the-art for $k > 1$ is an **energy reduction argument**.

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.
- Method of Bhattachrya et al does not generalize.
  Moral: The jump from one to two dimensions is key.
- Instead, state-of-the-art for $k > 1$ is an **energy reduction argument**.
  The **energy** of a Hegselmann-Krause system $\mathbf{x} = (x_1, \dots, x_n)$ is given by

$$\mathcal{E}(\mathbf{x}) = \sum_{i,j=1}^{n} \max\{1, ||x_i - x_j||^2\}.$$

Convergence in $\mathbb{R}^k$, $k > 1$:

- Opinions still freeze in finite time: this just requires a "convexity argument", which works in any Euclidean space.
- Method of Bhattachrya et al does not generalize.
  Moral: The jump from one to two dimensions is key.
- Instead, state-of-the-art for $k > 1$ is an **energy reduction argument**.
  The **energy** of a Hegselmann-Krause system $\mathbf{x} = (x_1, \dots, x_n)$ is given by

$$\mathcal{E}(\mathbf{x}) = \sum_{i,j=1}^{n} \max\{1, ||x_i - x_j||^2\}.$$

Basic Result: The dynamics always decrease the energy.

$$\mathcal{E}(\mathbf{x}(t)) - \mathcal{E}(\mathbf{x}(t+1)) \geq 4 \cdot \sum_{i=1}^{n} ||x_i(t) - x_i(t+1)||^2.$$

Convergence in $\mathbb{R}^k$, continued:

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.
  His result is in a sense best-possible. No better bound can be obtained using just the energy reduction technique.

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.
  His result is in a sense best-possible. No better bound can be obtained using just the energy reduction technique.
- **N.B.** The above only works for the L^2-norm.

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.
  His result is in a sense best-possible. No better bound can be obtained using just the energy reduction technique.
- **N.B.** The above only works for the L^2-norm.
- For lower bounds, n agents placed equidistantly around a circle will also require time $\Omega(n^2)$ to freeze.

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.
  His result is in a sense best-possible. No better bound can be obtained using just the energy reduction technique.
- **N.B.** The above only works for the L^2-norm.
- For lower bounds, n agents placed equidistantly around a circle will also require time $\Omega(n^2)$ to freeze.
  This is a genuinely 2-dimensional example. Also, in contrast to the dumbbell, this configuration reaches consensus.

Convergence in $\mathbb{R}^k$, continued:

- Martinsson [M, 2015] proved a bound of $O(n^4)$, valid in all dimensions.
  His result is in a sense best-possible. No better bound can be obtained using just the energy reduction technique.
- **N.B.** The above only works for the L^2-norm.
- For lower bounds, n agents placed equidistantly around a circle will also require time $\Omega(n^2)$ to freeze.
  This is a genuinely 2-dimensional example. Also, in contrast to the dumbbell, this configuration reaches consensus.
- We believe that the freezing time is $O(n^2)$ in all dimensions.

Convergence on $\mathbb{T}^1$:

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.

  E.g.: agents spaced equally around the circle at distance one.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.

  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.
  Proven by us [HMW].

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.
  Proven by us [HMW]. Proof uses both the energy reduction technqiue and a modification of the idea in Bhattachrya et al, both suitably modified for the circle.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.
  Proven by us [HMW]. Proof uses both the energy reduction technqiue and a modification of the idea in Bhattachrya et al, both suitably modified for the circle.
- The influence digraph can change at most $O(n^4)$ times.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.
  Proven by us [HMW]. Proof uses both the energy reduction technqiue and a modification of the idea in Bhattachrya et al, both suitably modified for the circle.
- The influence digraph can change at most $O(n^4)$ times.
  However, it can take arbitrarily long for these changes to occur.

Convergence on $\mathbb{T}^1$:

- In contrast to the Euclidean case, configurations no longer need to freeze in finite time.
- Moreover, in a frozen configuration, no cluster need be isolated.
  E.g.: agents spaced equally around the circle at distance one.
  However, there are even non-periodic frozen configurations.
- Hendrickx et al (2009) asked if opinions must always converge on the circle.
  Proven by us [HMW]. Proof uses both the energy reduction technqiue and a modification of the idea in Bhattachrya et al, both suitably modified for the circle.
- The influence digraph can change at most $O(n^4)$ times.
  However, it can take arbitrarily long for these changes to occur.
- Can also prove convergence in $\mathbb{T}^k$ for all $k \geq 1$ (technical).

In the simplest case, initial opinions are drawn **independently** from some probability distribution with compact support.

In the simplest case, initial opinions are drawn **independently** from some probability distribution with compact support.

Two basic principles, if applicable here, would together lead us to expect typical **phase transition** behaviour.

In the simplest case, initial opinions are drawn **independently** from some probability distribution with compact support.

Two basic principles, if applicable here, would together lead us to expect typical **phase transition** behaviour.

- **Monotonicity:**

In the simplest case, initial opinions are drawn **independently** from some probability distribution with compact support.

Two basic principles, if applicable here, would together lead us to expect typical **phase transition** behaviour.

- **Monotonicity:** Dilating the opinion space without changing the "relative distribution" of opinions should always make consensus less likely.

In the simplest case, initial opinions are drawn **independently** from some probability distribution with compact support.

Two basic principles, if applicable here, would together lead us to expect typical **phase transition** behaviour.

- **Monotonicity:** Dilating the opinion space without changing the “relative distribution” of opinions should always make consensus less likely.
- **Zero-One Law:**

In the simplest case, initial opinions are drawn **independently** from some probability distribution with compact support.

Two basic principles, if applicable here, would together lead us to expect typical **phase transition** behaviour.

- **Monotonicity:** Dilating the opinion space without changing the "relative distribution" of opinions should always make consensus less likely.
- **Zero-One Law:** Suppose initial opinions are chosen *independently* from some fixed distribution with compact support. As $n \to \infty$, the probability of reaching consensus should go to 0 or 1, i.e.: there should be a "typical behaviour".

In the simplest case, initial opinions are drawn **independently** from some probability distribution with compact support.

Two basic principles, if applicable here, would together lead us to expect typical **phase transition** behaviour.

- **Monotonicity:** Dilating the opinion space without changing the "relative distribution" of opinions should always make consensus less likely.
- **Zero-One Law:** Suppose initial opinions are chosen *independently* from some fixed distribution with compact support. As $n \to \infty$, the probability of reaching consensus should go to 0 or 1, i.e.: there should be a "typical behaviour".

Nothing is yet proven. Indeed, evidence against monotonicity is the fact that increasing the confidence bound r can sometimes destroy consensus.

Simulations:

Simulations:

- Many simulations performed in [DEJK, 2015] for uniform distributions of agents in regions of $\mathbb{R}^1$ and $\mathbb{R}^2$.

Simulations:

- Many simulations performed in [DEJK, 2015] for uniform distributions of agents in regions of $\mathbb{R}^1$ and $\mathbb{R}^2$.
- In $\mathbb{R}^1$ there is only one "region", namely an interval.

Simulations:

- Many simulations performed in [DEJK, 2015] for uniform distributions of agents in regions of $\mathbb{R}^1$ and $\mathbb{R}^2$.
- In $\mathbb{R}^1$ there is only one "region", namely an interval. Simulations give evidence for existence of a critical length, slightly above 5.

Simulations:

- Many simulations performed in [DEJK, 2015] for uniform distributions of agents in regions of $\mathbb{R}^1$ and $\mathbb{R}^2$.
- In $\mathbb{R}^1$ there is only one "region", namely an interval. Simulations give evidence for existence of a critical length, slightly above 5.

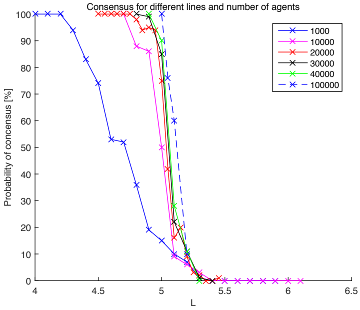

In any case, one approach is to study exhaustively a simpler class of configurations and then develop some approximation/interpolation theory.

In any case, one approach is to study exhaustively a simpler class of configurations and then develop some approximation/interpolation theory.

There exist two basic ideas here:

In any case, one approach is to study exhaustively a simpler class of configurations and then develop some approximation/interpolation theory.

There exist two basic ideas here:

Idea 1: Go to the limit of a continuum of agents.

In any case, one approach is to study exhaustively a simpler class of configurations and then develop some approximation/interpolation theory.

There exist two basic ideas here:

Idea 1: Go to the limit of a continuum of agents.

Idea 2: Study configurations of equally spaced agents.

Idea 1: The Continuous Agent Model (CAM)

Idea 1: The Continuous Agent Model (CAM)

- **Basic Idea:** Instead of drawing opinions independently from a (continuous) distribution $f(x)$, consider a continuum of agents with $f(x)$ describing an opinion density function.

Idea 1: The Continuous Agent Model (CAM)

- **Basic Idea:** Instead of drawing opinions independently from a (continuous) distribution $f(x)$, consider a continuum of agents with $f(x)$ describing an opinion density function.
- **Simplest example:** Uniformly independent opinions on $[0, L]$ corresponds to the opinion function $x_0 : [0, 1] \to \mathbb{R}$, $x_0(\alpha) = L\alpha$.

Idea 1: The Continuous Agent Model (CAM)

- **Basic Idea:** Instead of drawing opinions independently from a (continuous) distribution $f(x)$, consider a continuum of agents with $f(x)$ describing an opinion density function.
- **Simplest example:** Uniformly independent opinions on $[0, L]$ corresponds to the opinion function $x_0 : [0, 1] \to \mathbb{R}$, $x_0(\alpha) = L\alpha$.
- **The dynamics:**

$$x_{t+1}(\alpha) = \frac{1}{\mu(\mathcal{N}_t(\alpha))} \int_{\mathcal{N}_t(\alpha)} x_t(\beta)\, d\beta,$$

where $\mathcal{N}_t(\alpha) = \{\beta : ||x_t(\beta) - x_t(\alpha)|| \leq 1\}$ and μ is Lebesgue measure.

Idea 1: The Continuous Agent Model (CAM)

- **Basic Idea:** Instead of drawing opinions independently from a (continuous) distribution $f(x)$, consider a continuum of agents with $f(x)$ describing an opinion density function.
- **Simplest example:** Uniformly independent opinions on $[0, L]$ corresponds to the opinion function $x_0 : [0, 1] \to \mathbb{R}$, $x_0(\alpha) = L\alpha$.
- **The dynamics:**

$$x_{t+1}(\alpha) = \frac{1}{\mu(\mathcal{N}_t(\alpha))} \int_{\mathcal{N}_t(\alpha)} x_t(\beta)\, d\beta,$$

  where $\mathcal{N}_t(\alpha) = \{\beta : ||x_t(\beta) - x_t(\alpha)|| \leq 1\}$ and μ is Lebesgue measure.
- **Approximation between DAM and CAM:** Hendrickx et al (2009) have results which are probably strong enough for most purposes.

Idea 1: The Continuous Agent Model (CAM)

- **Basic Idea:** Instead of drawing opinions independently from a (continuous) distribution $f(x)$, consider a continuum of agents with $f(x)$ describing an opinion density function.
- **Simplest example:** Uniformly independent opinions on $[0, L]$ corresponds to the opinion function $x_0 : [0, 1] \to \mathbb{R}$, $x_0(\alpha) = L\alpha$.
- **The dynamics:**

$$x_{t+1}(\alpha) = \frac{1}{\mu(\mathcal{N}_t(\alpha))} \int_{\mathcal{N}_t(\alpha)} x_t(\beta)\, d\beta,$$

  where $\mathcal{N}_t(\alpha) = \{\beta : ||x_t(\beta) - x_t(\alpha)|| \leq 1\}$ and μ is Lebesgue measure.
- **Approximation between DAM and CAM:** Hendrickx et al (2009) have results which are probably strong enough for most purposes. So it remains to study the CAM-model.

- However, this doesn't seem to be at all straightforward. Not clear if we're really simplifying things with CAM.

- However, this doesn't seem to be at all straightforward. Not clear if we're really simplifying things with CAM.
- It is not even known in general if a configuration of opinions in CAM always converges to something (Hendrickx et al, 2009).

- However, this doesn't seem to be at all straightforward. Not clear if we're really simplifying things with CAM.
- It is not even known in general if a configuration of opinions in CAM always converges to something (Hendrickx et al, 2009).
- Wedin and I [WH1] gave the first example of a **regular** opinion function (piecewise differentiable, with positive lower and upper bounds on the derivative) which never reaches consensus.

- However, this doesn't seem to be at all straightforward. Not clear if we're really simplifying things with CAM.
- It is not even known in general if a configuration of opinions in CAM always converges to something (Hendrickx et al, 2009).
- Wedin and I [WH1] gave the first example of a **regular** opinion function (piecewise differentiable, with positive lower and upper bounds on the derivative) which never reaches consensus. Even this is a non-trivial task.

- However, this doesn't seem to be at all straightforward. Not clear if we're really simplifying things with CAM.
- It is not even known in general if a configuration of opinions in CAM always converges to something (Hendrickx et al, 2009).
- Wedin and I [WH1] gave the first example of a **regular** opinion function (piecewise differentiable, with positive lower and upper bounds on the derivative) which never reaches consensus. Even this is a non-trivial task.
- Our example is a kind of **double-S**.

- However, this doesn't seem to be at all straightforward. Not clear if we're really simplifying things with CAM.
- It is not even known in general if a configuration of opinions in CAM always converges to something (Hendrickx et al, 2009).
- Wedin and I [WH1] gave the first example of a **regular** opinion function (piecewise differentiable, with positive lower and upper bounds on the derivative) which never reaches consensus. Even this is a non-trivial task.
- Our example is a kind of **double-S**.
- Problem remains open for linear functions (those corresponding to a uniform distribution of opinions).

Idea 2: Equally spaced agents

Idea 2: Equally spaced agents

- Recall that in [HW] we proved that a configuration of n agents, with initial opinions $0, 1, \ldots, n-1$, evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.

Idea 2: Equally spaced agents

- Recall that in [HW] we proved that a configuration of n agents, with initial opinions $0, 1, \ldots, n-1$, evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.
- It is conceptually easier to consider a semi-infinite configuration of equally spaced agents, with initial opinions at all non-negative integers.

Idea 2: Equally spaced agents

- Recall that in [HW] we proved that a configuration of n agents, with initial opinions $0, 1, \ldots, n-1$, evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.
- It is conceptually easier to consider a semi-infinite configuration of equally spaced agents, with initial opinions at all non-negative integers.
  The first (and main) step in [HW] was to prove that this configuration evolves periodically, with a group of 3 agents breaking loose on the left after every 5th time step.

Idea 2: Equally spaced agents

- Recall that in [HW] we proved that a configuration of n agents, with initial opinions $0, 1, \dots, n-1$, evolves periodically, with groups of 3 agents breaking loose at each end every 5th time step.
- It is conceptually easier to consider a semi-infinite configuration of equally spaced agents, with initial opinions at all non-negative integers.
  The first (and main) step in [HW] was to prove that this configuration evolves periodically, with a group of 3 agents breaking loose on the left after every 5th time step.
- Now one should consider a general inter-agent spacing $d \in (0, 1]$ - ultimately we are interested in letting $d \to 0$.

- The previously described behaviour actually occurs for any $d > 72/79$.

- The previously described behaviour actually occurs for any $d > 72/79$.
- We don't know any value of d where the evolution does not appear to be *ultimately* periodic.

- The previously described behaviour actually occurs for any $d > 72/79$.
- We don't know any value of d where the evolution does not appear to be *ultimately* periodic. We can prove that this is always so when $d > 1/2$,

- The previously described behaviour actually occurs for any $d > 72/79$.
- We don't know any value of d where the evolution does not appear to be *ultimately* periodic. We can prove that this is always so when $d > 1/2$, where basically only 12 different "kinds of behaviour" are possible, though a system may jump from one kind to another before settling down (hence an ultimately periodic, but not periodic evolution).

- The previously described behaviour actually occurs for any $d > 72/79$.
- We don't know any value of d where the evolution does not appear to be *ultimately* periodic. We can prove that this is always so when $d > 1/2$, where basically only 12 different "kinds of behaviour" are possible, though a system may jump from one kind to another before settling down (hence an ultimately periodic, but not periodic evolution).
- Wedin is working on developing an appropriate approximation/interpolation theory.

- The previously described behaviour actually occurs for any $d > 72/79$.
- We don't know any value of d where the evolution does not appear to be *ultimately* periodic. We can prove that this is always so when $d > 1/2$, where basically only 12 different "kinds of behaviour" are possible, though a system may jump from one kind to another before settling down (hence an ultimately periodic, but not periodic evolution).
- Wedin is working on developing an appropriate approximation/interpolation theory.
- Most intriguingly, simulations suggest a possible triple phase transition !